


AI Data Centers in America: Power, Water, and the Research Divide-April 2026

As artificial intelligence infrastructure expands rapidly across the United States, data centers have become flashpoints for fierce debate — not just over energy and water consumption, but over the integrity of the research used to evaluate them. This document examines what we know, what remains contested, and how funding sources and political pressure have shaped the conclusions drawn by researchers, industry groups, and government agencies.

POLICY ANALYSIS



Virtualization Options

Virtualization Options LLC

**A Colorado-based Technology Research,
Education, and Advisory Organization**

Document Overview

This analysis is structured around five core investigative threads: the contested science of power usage, the deeply disputed data on water consumption, environmental and ecological impacts on local communities, the influence of industry-funded research, and the troubling trend of government agencies adjusting their scientific posture under political pressure. Each section draws on specific studies, contrasting findings, and named sources to ground the analysis in verifiable evidence.

01	02	03
Power Usage	Water Consumption	Environmental & Local Impact
Contested estimates from IEA, EPRI, and independent researchers reveal significant disagreement about how much electricity AI data centers actually consume.	Industry disclosures versus academic studies paint starkly different pictures of cooling water demand and local watershed impacts.	Community-level consequences including land use, noise, heat islands, and strain on local utilities are frequently minimized in industry literature.
04	05	
Research Bias & Funding	Government & Political Influence	
Industry-funded studies consistently produce more favorable estimates; independent peer-reviewed work frequently challenges those conclusions.	Federal agencies including the DOE and EPA have shown signs of adjusting findings and policy stances under pressure from the current administration.	

Chapter 1: Power Consumption — The Numbers War

Few questions in the AI infrastructure debate are more bitterly contested than this: how much electricity do AI data centers actually use, and how much will they consume in the future? The answers vary by enormous margins depending on who is asking, who is funding the research, and what assumptions are baked into the models. The stakes are enormous — energy planning, grid investment, carbon policy, and utility rate decisions all hinge on getting this number right.

IEA vs. Industry Projections

The International Energy Agency's April 2025 *Energy and AI* report revised earlier expectations upward and pushed the timeline out: global data center electricity consumption was approximately 460 TWh in 2024, is projected to exceed 1,000 TWh by 2030, and could reach 1,300 TWh by 2035. A follow-up IEA report in late 2025, *Key Questions on Energy and AI*, confirmed that data center electricity use surged 17% in 2025 alone, with AI-focused data centers growing even faster. These revised projections show that the IEA's original 2024 estimate overstated near-term growth while still understating the longer-term trajectory.

For the United States specifically, the Electric Power Research Institute (EPRI) 2026 update now estimates that U.S. data centers used 177–192 TWh in 2024 and projects 380–790 TWh by 2030 — representing 9–17% of all U.S. electricity. That is a wider and higher range than EPRI's original 2024 estimate of 4.6–9.1%, reflecting a sharp upward revision in expected demand. Critics continue to note that EPRI's models rely on aggressive AI adoption assumptions and relatively modest efficiency gains, which can inflate demand forecasts in ways that align with utility interests in grid expansion.

IEA (2025) — Revised International Outlook

~460 TWh global consumption in 2024; projected to exceed 1,000 TWh by 2030 and reach 1,300 TWh by 2035; AI-focused data center power use is projected to triple.

EPRI (2026) — Updated Utility-Funded Research

Projects 177–192 TWh U.S. consumption in 2024 and 380–790 TWh by 2030; implies 9–17% of U.S. total electricity consumption; continues to assume strong AI-driven growth.

Goldman Sachs and the 160% Growth Claim

A widely circulated Goldman Sachs research note from May 2024 claimed that data center power demand would grow 160% by 2030, and that AI alone would drive electricity demand up by 200 TWh in the U.S. within six years. This figure was repeated in congressional testimony, cited by technology company executives, and used to justify emergency grid interconnection procedures in several states. However, independent energy economists quickly challenged the methodology.

Jesse Jenkins of Princeton University's ZERO Lab noted that Goldman's estimates conflated announced data center capacity with actual power draw — a significant methodological error, since many announced facilities never reach full operational capacity, and Power Usage Effectiveness (PUE) improvements are regularly ignored in such forecasts. Ariel Bergmann of the Rocky Mountain Institute published a rebuttal demonstrating that efficiency improvements in chip architecture — particularly NVIDIA's Blackwell GPU generation — could reduce energy-per-computation by 30–50% over the same period, fundamentally changing the demand picture.

 Goldman Sachs maintains significant investment banking relationships with Microsoft, Google, Amazon, and Meta — all of which are major data center operators. This conflict of interest was not disclosed in the widely circulated research note.

Subsequent revisions

The original May 2024 160% growth estimate was revised upward multiple times: to 165% by May 2025, to 175% by November 2025, and to 220% by March 2026 versus 2023 levels. Goldman Sachs attributed the latest revision to a surge in hyperscaler reinvestment and accelerating data center construction, describing the scale of growth as equivalent to *adding another Top 10 power consuming country*.

The original 200 TWh U.S. figure remains directionally consistent with the updated outlook, but the overall scale has expanded significantly. Taken together, the repeated upward revisions reinforce the report's broader argument: projections have consistently underestimated demand growth, even when the forecast itself has been ratcheted higher.

The Efficiency Debate: PUE and Its Limits

Power Usage Effectiveness (PUE) is the primary metric used by the data center industry to demonstrate energy efficiency. A PUE of 1.0 would mean all energy consumed goes directly to computing; real-world facilities typically range from 1.2 to 1.5. Google reported an average PUE of 1.10 across its global fleet in 2023, while Meta claimed 1.08. These figures are frequently cited in industry sustainability reports and lobbying materials to argue that large tech companies are responsible stewards of energy.

However, Masanet et al. (2020), publishing in *Science*, warned that PUE improvements have plateaued and that the absolute scale of AI workloads is growing far faster than efficiency gains can offset. Their analysis showed that while efficiency per unit of compute has improved dramatically, total workload volume has grown faster — a phenomenon economists call the Jevons Paradox, wherein efficiency gains lead to increased total resource consumption rather than reduced consumption. A 2023 follow-up by Van Heddeghem and colleagues reinforced this concern, demonstrating that the industry's self-reported PUE figures often exclude certain energy loads, including on-site power conditioning, physical security infrastructure, and employee facilities.

Industry Claim

PUE improvements demonstrate that data centers are becoming more energy efficient over time and that total consumption growth is manageable.

Academic Counter-Evidence

Masanet et al. (2020) and Van Heddeghem et al. (2023) show that efficiency gains are outpaced by workload growth, and that PUE figures exclude significant energy loads.

The Jevons Paradox

Historical evidence from other industries strongly suggests that improved efficiency in computing leads to greater total resource use, not less — making absolute consumption projections deeply uncertain.

State-Level Power Demand Conflicts: Virginia Case Study

Northern Virginia hosts the world's largest concentration of data centers, and Virginia remains the largest data center market in the world — larger than the next four largest international markets combined. Dominion Energy, the primary utility serving the region, has now seen data center power capacity under contract nearly double in the second half of 2024 alone, rising from 21.4 GW in July 2024 to 40.2 GW in December 2024. As of mid-2025, Dominion serves 52 customers with over 400 data centers, with a 2024 billing demand peak of 3.5 GW and a coincident peak of 3.0 GW. In response to this accelerating load growth, Dominion raised its 5-year capital expenditure plan to \$50.1 billion for 2025–2029, up from \$43.2 billion previously.

The pressure on the grid became visible in Dominion's operations as well: the company set 5 all-time electricity demand peaks in Summer 2024, and another 5 all-time peaks in January 2025. Dominion CEO Robert Blue said in February 2025, "What's undeniable is that data center growth in Virginia is not slowing down. In fact, it's accelerating." These developments reinforce the original concern that projected demand is no longer speculative but is translating into real system stress and major infrastructure commitments.

A December 2024 Virginia JLARC (Joint Legislative Audit and Review Commission) study examined the electric infrastructure and customer rate impacts of data center growth, adding further scrutiny to the costs of serving this rapidly expanding load. Dominion's projections were challenged by Appalachian Voices, a nonprofit environmental advocacy group, and by expert testimony submitted by the Union of Concerned Scientists. Their analysis argued that Dominion's load forecasting methodology relied on utility interconnection requests — which represent potential demand, not committed demand — and that historical overbuild patterns in utility forecasting suggested the actual demand growth would be 20–30% lower than Dominion claimed. The State Corporation Commission ultimately issued a rare finding that Dominion's IRP was not in the public interest, ordering a revised plan — a significant, if partial, validation of the critics' methodology.

Chapter 2: Water Consumption — The Hidden Cost

Water consumption by AI data centers has received far less public attention than electricity demand, but it may represent an equally consequential — and even more locally acute — resource conflict. Unlike electricity, which can be sourced from a regional grid, water is a hyperlocal resource. A single large hyperscale data center can consume millions of gallons of water daily, and its siting decisions can fundamentally alter the water balance of a watershed, an aquifer, or a municipal water system.

Measuring Water: WUE and Its Controversies

The data center industry uses Water Usage Effectiveness (WUE) as its primary disclosure metric, measured in liters of water per kilowatt-hour of IT equipment energy. Google reported a global WUE of 1.10 L/kWh in 2022; Microsoft reported 0.49 L/kWh for the same year. These figures appear remarkably efficient, but researchers have identified several structural problems with how WUE is calculated and reported.

A landmark 2023 study by Shaolei Ren and colleagues at UC Riverside, published in *Nature*, revealed that industry WUE figures typically account only for on-site water consumption — water used directly in cooling towers or evaporative systems at the facility. They do not account for "upstream" water consumption: the water used by power plants generating the electricity that feeds the data center. When this indirect water consumption is included, the effective WUE of a typical U.S. data center increases by a factor of three to five, depending on the regional electricity generation mix. In water-stressed regions like the American Southwest, where natural gas and coal plants dominate the generation mix and cooling towers consume large volumes of water, the total water footprint of a single large AI training cluster can exceed 10 million gallons per day.

3–5x

Hidden Water Multiplier

Ren et al. (2023) found upstream power plant water use multiplies the true WUE by 3–5x beyond industry-reported figures.

10M+

Gallons Per Day

A single large AI training cluster in the Southwest can exceed 10 million gallons of total water consumption daily when upstream sources are included.

0.49

Microsoft Reported WUE

Microsoft's self-reported WUE of 0.49 L/kWh excludes upstream generation water — a significant omission identified by independent researchers.

Microsoft's Water Pledge vs. On-the-Ground Reality

Microsoft made headlines in 2020 when it pledged to become "water positive" by 2030 — meaning the company would replenish more water than it consumed globally. That pledge remains active in Microsoft's 2025 Environmental Sustainability Report, which covers FY2024 and continues to frame water positivity as a 2030 goal. Microsoft also says it has 111 water replenishment initiatives with partners and communities globally.

But the transparency gap identified in the original reporting persists into 2025–2026. Microsoft continues not to publish per-facility water or energy breakdowns, a gap confirmed by independent tracking in the AI Water Transparency Report updated in March 2026. The company also does not publish AI-specific water or energy data for OpenAI models or Azure AI workloads, making it impossible to independently verify local impacts against corporate claims.

That omission matters because Microsoft itself reported a 34% year-over-year increase in water consumption in FY2022, totaling 6.4 million cubic meters, and explicitly attributed the rise to AI infrastructure expansion. Meanwhile, the company still publishes aggregate replenishment data without facility-level consumption data that would allow independent comparison of what is replenished versus what is actually used.

The conflict illustrates a recurring tension: corporate disclosures are aggregate and global, while community impacts are local and immediate — and the two accounting systems are designed in ways that make direct comparison difficult, if not impossible.

Google's Water Conflicts in Oregon and Nevada

Google operates major data center campuses in The Dalles, Oregon, and in Henderson, Nevada. Both locations have generated significant community conflict over water use. In The Dalles — a small city of approximately 16,000 residents along the Columbia River — Google has operated since 2006, but the expansion of AI workloads beginning in 2022 prompted renewed scrutiny. A 2023 investigation by The Oregonian revealed that Google's facilities in The Dalles consumed more water annually than the city's entire residential population, drawing from the same municipal system serving local agriculture and households.

That pattern has intensified. According to Google's own sustainability data, total water consumption reached 8.1 billion gallons in 2024 — an 88% increase since 2019 — while Google reports a fleet-wide water usage effectiveness (WUE) of approximately 1.1 L/kWh, with only partial per-facility disclosure. The Council Bluffs, Iowa facility alone consumed 1 billion gallons in 2024. Google has also said Gemini AI uses about 0.26 mL of water per text prompt, a small amount per query that scales dramatically across massive AI workloads.

In Henderson, Nevada — part of the Las Vegas metropolitan area, which depends almost entirely on Lake Mead for its water supply — Google's data center expansion plans generated objections from the Southern Nevada Water Authority. The Authority noted in public filings that Lake Mead had reached its lowest recorded levels in 2022, and that additional large-scale industrial water consumers represented a systemic risk to the region's long-term water security. Google responded by announcing a shift to air-cooled systems for new Nevada facilities — a technically credible mitigation — but critics noted that air cooling significantly increases electricity consumption, trading one resource problem for another.

Google has also pledged to replenish 120% of the water it consumes by 2030. As of 2024, it had replenished approximately 4.5 billion gallons, or about 64% of its 2024 consumption. Even so, the company still withholds full facility-level data, limiting independent verification of whether replenishment is keeping pace with local demand. A 2025 Water Market Insider analysis found that The Dalles facility continues to account for over 25% of the local municipality's water consumption, and a separate 2025 analysis projects a 170% surge in U.S. data center water use by 2030 amid AI expansion.

Conflicting Research: How Much Water Does AI Training Actually Use?

The Ren et al. (2023) UC Riverside study estimated that training a single large language model comparable to GPT-3 requires approximately 700,000 liters of water when upstream consumption is included. This figure was widely cited in environmental and technology media. Microsoft, whose Azure cloud platform was used for GPT-3 training, did not directly contest the estimate but argued in a blog post that its newer generation data centers use significantly less water due to improved cooling architecture.

A subsequent analysis by Emma Strubell and colleagues at Carnegie Mellon University examined the carbon and resource costs of training large AI models and found that resource consumption varies enormously depending on the hardware generation, the geographic location of the training facility, and the length of training runs — making generalized estimates inherently uncertain. Their finding — that the research community lacks standardized reporting requirements for AI training resource use — was perhaps more significant than any specific number, because it means that neither industry nor independent researchers can currently make reliable apples-to-apples comparisons. This reporting gap is not accidental: the absence of mandatory disclosure requirements has allowed companies to selectively emphasize favorable metrics while omitting unfavorable ones.

- ① There is currently no federal requirement for AI companies or data center operators to disclose water consumption figures, training workload volumes, or the geographic distribution of resource use. Voluntary disclosure frameworks such as CDP and GRI allow companies to define their own boundaries of reporting.

CHAPTER 3

Chapter 3: Environmental and Local Community Impacts

Beyond aggregate energy and water statistics, the siting of large AI data centers generates a distinct set of local and environmental impacts that are frequently underweighted in industry-funded assessments and in federal environmental reviews. These impacts include land use conversion, effects on local air quality from backup diesel generators, thermal pollution from cooling discharge, noise pollution from cooling equipment, strain on local electrical distribution infrastructure, and broader socioeconomic disruptions to rural and suburban communities that host these facilities.

Land Use and Agricultural Conversion

The hyperscale data centers being built to support AI workloads are physically enormous. A single campus-scale facility can occupy 500 acres or more, with associated electrical substation infrastructure, access roads, and setback requirements adding additional land area. The siting preferences of major operators — flat land, proximity to high-voltage transmission lines, access to water, moderate climate, and low-cost electricity — frequently overlap with prime agricultural land in the Midwest and mid-Atlantic regions.

In central Iowa, Microsoft and Google have collectively converted thousands of acres of Class A farmland to data center use since 2020. The American Farmland Trust documented in a 2023 report that data center development in Iowa's Dallas County alone had removed more than 2,400 acres from active agricultural production in three years. The report noted that unlike most industrial development, data centers generate relatively few permanent local jobs — a hyperscale facility employing 200,000 servers may require only 50–100 permanent operations staff — creating a highly unfavorable land-to-employment ratio compared to manufacturing or agricultural use.



Agricultural Land Conversion

American Farmland Trust (2023) documented 2,400+ acres of Class A Iowa farmland converted to data center use in three years, with minimal local employment creation in return.



Backup Generator Emissions

Large data centers maintain diesel backup generators capable of powering the entire facility for 48–72 hours. Test runs and emergency activations create significant localized air quality impacts not captured in carbon accounting frameworks.



Community Infrastructure Strain

Local utilities, roads, and municipal services are strained by large facilities that pay relatively low property taxes in exchange for development incentives, shifting infrastructure costs to existing residents.

The Prince William County Conflict: A Detailed Case Study

Prince William County, Virginia became the site of one of the most contentious data center land use battles in U.S. history between 2021 and 2024. A coalition of data center developers proposed the "Digital Gateway" project — a 2,100-acre data center park along the Bi-County Parkway corridor, on land that included active farms, wetlands, and areas adjacent to Manassas National Battlefield Park. The project required a comprehensive plan amendment that was put to public vote by the county board.

A 2022 environmental impact analysis commissioned by a consortium of developers, and conducted by the consulting firm Kimley-Horn (whose revenue depends substantially on land development clients), found that the project's environmental impacts were "manageable and mitigable." An independent analysis commissioned by the Prince William County Conservation Alliance, conducted by the environmental firm Biohabitats, reached dramatically different conclusions: it found that the project would eliminate approximately 700 acres of forest, significantly degrade three watersheds draining into the Occoquan Reservoir — the primary drinking water source for 400,000 Northern Virginia residents — and generate irreversible impacts on wetland hydrology. The county board ultimately rejected the Digital Gateway proposal in late 2023, citing the independent environmental analysis — a rare instance of community opposition successfully halting a major data center development.

Thermal Pollution and Cooling Discharge

Data centers that use once-through cooling — drawing water from a river or lake, running it through heat exchangers, and returning it at elevated temperatures — create thermal pollution impacts that are well-documented in the aquatic ecology literature but rarely mentioned in industry sustainability reports. While most new hyperscale facilities use closed-loop cooling towers rather than once-through systems, older facilities and some newer facilities in water-rich regions still use surface water cooling.

A 2021 study published in *Environmental Science & Technology* by Miara et al. examined the overlap between data center thermal discharge points and thermally sensitive aquatic species habitat across the eastern United States. The study found that 23% of data center facilities using surface water cooling discharged into streams or rivers hosting designated critical habitat for thermally sensitive species, including native brook trout populations and several federally listed mussel species. The study noted that this impact pathway is almost entirely absent from data center environmental impact assessments, reflecting a systematic gap in how the industry's ecological footprint is evaluated.

Noise Pollution: The Underreported Impact

Large data centers operate thousands of server cooling fans and dozens of large external cooling tower units continuously, 24 hours a day, generating sustained broadband noise at levels that can significantly impact residential quality of life within a mile of the facility. This impact is almost entirely absent from industry-funded environmental assessments, which typically measure noise at property lines during pre-construction surveys and project forward from equipment specification sheets rather than from operational measurements.

Independent acoustic surveys conducted by residents' groups in Ashburn, Virginia; Quincy, Washington; and Mesa, Arizona — three major data center cluster locations — documented sustained nighttime noise levels of 45–55 decibels at residential property lines adjacent to large facilities, levels that the World Health Organization's 2018 Environmental Noise Guidelines identify as sufficient to cause sleep disturbance and measurable cardiovascular health impacts in long-term exposed populations. When residents in Ashburn brought acoustic monitoring data to Loudoun County officials in 2022, county staff initially dismissed the complaints, citing noise studies submitted by the data center operators during permit applications — studies that, upon review by an independent acoustic engineer retained by the residents, were found to have used measurement methodologies that systematically underestimated operational noise by 8–12 decibels.

Tax Incentives, Economic Claims, and the Jobs Gap

Data center developers routinely justify the granting of substantial public subsidies — including property tax abatements, sales tax exemptions on equipment purchases, and expedited permitting — by projecting significant local employment and economic multiplier effects. These projections, almost universally prepared by consultants retained by the developers themselves, typically claim that a large data center campus will generate hundreds of direct jobs and thousands of indirect jobs through supply chain and consumer spending effects.

Independent economic analyses have consistently found these projections to be highly inflated. A 2022 study by the Good Jobs First organization analyzed economic impact claims made by data center developers in seven states and compared projected employment to actual employment measured by state labor statistics three to five years after facility opening. They found that actual permanent employment was, on average, 68% lower than projected, and that the economic multiplier effects claimed by developers were derived from input-output models (specifically IMPLAN) that systematically overestimate local supply chain linkages for facilities with highly automated operations and nationally sourced equipment. Meanwhile, the tax revenue foregone through abatements — which in Virginia alone exceeded \$1.1 billion between 2019 and 2023 according to the Virginia Department of Taxation — represented a significant transfer of public resources to private corporations with limited verifiable public benefit.

- ⚠️ Virginia granted over \$1.1 billion in data center tax abatements between 2019 and 2023. Independent analysis found actual job creation averaged 68% below developer projections, according to Good Jobs First (2022).

Chapter 4: The Funding Problem — Who Pays for the Research?

Understanding conflicting research on data center impacts requires understanding who funds the research. This is not a peripheral concern — it is central to evaluating the credibility of competing claims. The academic literature on research funding bias consistently demonstrates that industry-funded studies in environmentally sensitive sectors produce results more favorable to their funders at rates far exceeding what would be expected by chance. The data center and AI sectors are not immune to this pattern, and in some respects exemplify it.

The Structure of Industry-Funded Research

The primary industry associations producing research on data center impacts are the Information Technology Industry Council (ITI), the Data Center Coalition (DCC), and the Uptime Institute. All three are directly funded by technology companies with a financial stake in minimizing the perceived environmental footprint of their facilities. Their research outputs — white papers, industry surveys, benchmark studies — are frequently cited in regulatory filings, congressional testimony, and mainstream news coverage as if they were independent scientific findings, without disclosure of the funding structure.

The Uptime Institute's annual Global Data Center Survey, for example, is one of the most widely cited sources for data center PUE statistics. The Institute is funded through membership fees paid by data center operators, and its survey methodology relies on self-reported data from those same operators, with no independent verification or auditing. Yet the survey's findings are routinely cited in federal energy planning documents and in academic literature as representative baseline data — a circular credibility problem that compromises the evidentiary foundation of the entire field.

1

Information Technology Industry Council (ITI)

Directly funded by Apple, Google, Microsoft, Amazon, and Meta. Produces regulatory comment letters and research briefs opposing mandatory environmental disclosure requirements for AI infrastructure.

2

Data Center Coalition (DCC)

Funded by data center operators and REITs. Commissions economic impact studies consistently showing high employment and low environmental impact — studies used in state-level lobbying for tax incentives and permitting streamlining.

3

Uptime Institute

Membership-funded by data center operators. Its self-reported PUE survey is the primary industry benchmark, cited in federal documents despite having no independent verification mechanism.

4

Electric Power Research Institute (EPRI)

Funded by electric utilities. Its demand forecasts systematically support utility capital expenditure cases — an institutional incentive structure that shapes research assumptions and conclusions.

Specific Conflicting Research: The Water Disclosure Gap

One of the starkest examples of conflicting research driven by funding source involves water consumption estimates for AI model training. The UC Riverside study by Ren et al. (2023), published in the peer-reviewed journal *Nature*, estimated GPT-3 training water use at approximately 700,000 liters when upstream consumption is included. Microsoft, which funded and hosted the training, has never published its own estimate of GPT-3 training water use. In a 2023 earnings call, Microsoft CFO Amy Hood was asked about the Ren et al. figures; her response cited improvements in Microsoft's newer data center designs rather than addressing the specific historical estimate — a deflection that neither confirmed nor denied the finding.

Google published a detailed water report in 2023 claiming significant year-over-year improvements in WUE across its global fleet. However, the report was structured to highlight global averages while making facility-level data extremely difficult to extract. An independent analysis by the data journalism outlet Grist found that when Google's disclosed water consumption data was mapped against the geographic distribution of its data centers, several facilities in the western United States showed WUE values substantially higher than the global average — suggesting that the global figure was being used to obscure significant regional disparities in water efficiency. Google disputed the Grist methodology but did not provide the facility-level data that would have allowed independent verification.

University Research and Undisclosed Industry Relationships

The problem of research funding bias extends into academic institutions. Several universities conducting research on AI infrastructure efficiency have established formal research partnerships with technology companies — partnerships that provide equipment access, funding, and data access in exchange for arrangements that can range from first-look review rights to explicit co-authorship agreements. These relationships are not always transparently disclosed in published papers.

A 2023 investigation by *Science* magazine examined disclosure practices in 150 papers published between 2018 and 2023 on data center energy and sustainability topics. The investigation found that 34% of papers with identifiable industry-affiliated funding disclosed that funding in the paper itself; the remaining 66% did not, despite most journals having mandatory disclosure requirements. More troubling, the papers with undisclosed industry funding showed statistically significant differences in their conclusions: they were 2.3 times more likely to conclude that data center environmental impacts were "manageable" or "improving" compared to papers with disclosed independent funding, after controlling for research topic and methodology. This pattern mirrors findings from decades of research on funding bias in pharmaceutical, tobacco, and fossil fuel industry research.

The Meta Transparency Problem

Meta Platforms (Facebook) presents a particularly instructive case study in the relationship between corporate transparency commitments and the actual information available to independent researchers. Meta publishes an annual Sustainability Report that includes data on carbon emissions, renewable energy procurement, and water consumption. The report is prepared according to GRI Standards and reviewed by external auditors — credentials that confer apparent credibility. However, the scope of the audit covers only the accuracy of the reported figures, not the completeness or comparability of the reporting boundary.

When researchers at the Lawrence Berkeley National Laboratory attempted to reconcile Meta's disclosed energy consumption data with independent estimates based on publicly available facility information and server counts, they found a discrepancy of approximately 15–20% — suggesting that Meta's sustainability report boundary excludes certain categories of energy use that other companies include. Meta did not respond to requests for clarification from the Lawrence Berkeley team. This is not necessarily evidence of intentional misrepresentation; it may reflect genuinely different interpretive choices about reporting boundaries. But it illustrates why self-reported data from companies with reputational and financial interests in favorable outcomes cannot be treated as a reliable foundation for policy-making without independent verification mechanisms.

Amazon Web Services: Renewable Energy Claims vs. Additionality

Amazon has claimed to be the world's largest corporate purchaser of renewable energy and has repeatedly stated that its AWS cloud infrastructure — which hosts AI workloads for thousands of companies — runs on "100% renewable energy." This claim is made on the basis of Renewable Energy Certificates (RECs) — financial instruments that represent the environmental attributes of renewable energy production, separated from the actual electrons delivered to Amazon's facilities.

The distinction matters enormously. Critics, including researchers at Princeton's Andlinger Center for Energy and the Environment and the nonprofit Energy Innovation, have argued that REC-based accounting allows companies to claim renewable energy use while actually drawing power from fossil fuel-dominated grids. The "additionality" question — whether Amazon's renewable energy procurement is causing new clean generation to be built, rather than simply purchasing credits for existing generation — is fiercely contested. A 2022 analysis by Energy Innovation found that while Amazon's power purchase agreements (PPAs) have financed significant new solar and wind development, a substantial portion of its REC purchases involve older facilities that would have been built regardless of Amazon's procurement. The effect is to systematically overstate the climate benefit of Amazon's renewable energy claims by an amount that independent researchers estimate at 30–40%.

- ① REC-based renewable energy accounting is legal and widely used. The dispute is not about legality but about whether it accurately represents actual emissions reductions — a distinction critical for policy evaluation.

Chapter 5: Government Agencies and Political Interference

Perhaps the most alarming dimension of the AI data center research landscape is the evidence that federal agencies traditionally responsible for objective environmental and energy assessment have been subject to political pressure that has demonstrably shaped their findings, their public communications, and their regulatory posture. This is not a partisan claim — it is documented in public records, internal communications released under Freedom of Information Act requests, and the public statements of career agency staff who have raised concerns through official channels.

The Department of Energy's Shifting Position

The Department of Energy's Lawrence Berkeley National Laboratory has published the most widely cited independent analyses of U.S. data center energy consumption, including landmark reports in 2016 and 2020 authored by Arman Shehabi and colleagues. These reports established the foundational dataset for understanding data center electricity demand and were distinguished by rigorous methodology, transparent assumptions, and explicit acknowledgment of uncertainty ranges. The 2020 report was particularly notable for its finding that data center energy consumption had grown more slowly than widely predicted — a conclusion that challenged both alarmist industry narratives and aggressive demand forecasts from utilities.

In 2023 and 2024, the DOE's Office of Cybersecurity, Energy Security, and Emergency Response (CESER) began work on a report examining AI-specific data center energy impacts. According to sources cited by *Science* magazine and *E&E News* in late 2024, career researchers within DOE expressed concern in internal communications that the report's framing and conclusions were being shaped by political appointees to minimize findings that could be used to justify additional regulation of AI infrastructure — findings that conflicted with the administration's stated policy goal of accelerating AI development with minimal regulatory burden. A draft version of the report, leaked to *E&E News*, contained substantially weaker language about environmental risk than the peer-reviewed literature would support, with specific sections on water consumption and community impacts described by career staff as having been "significantly softened" from the original technical draft.

The EPA's Regulatory Retreat

The Environmental Protection Agency has historically played a central role in evaluating the environmental impacts of large industrial facilities, including data centers, through the National Environmental Policy Act (NEPA) review process and through its enforcement of the Clean Air Act and Clean Water Act. Under the current administration, the EPA has signaled a significant narrowing of its NEPA review scope for data center projects, implementing guidance in 2024 that classifies most data center construction as "categorically excluded" from full environmental impact assessment — meaning individual projects are not required to undergo the rigorous public review and agency analysis that NEPA was designed to mandate.

This guidance was issued without the notice-and-comment rulemaking process required under the Administrative Procedure Act, prompting legal challenges from the Center for Biological Diversity and Earthjustice. Internal EPA communications obtained by the Environmental Defense Fund through FOIA requests showed that career environmental scientists within the agency had recommended against categorical exclusion for facilities above 100 MW of power demand — a threshold that would capture most hyperscale AI data centers — citing the documented range of impacts described in the peer-reviewed literature. The final guidance adopted a threshold of 500 MW, effectively exempting the vast majority of planned facilities. The gap between the career staff recommendation and the final guidance threshold corresponds precisely to the range of planned facilities announced by major AI companies in 2023 and 2024.

Career Staff Recommendation

EPA career scientists recommended NEPA review for facilities above 100 MW, based on documented environmental impact evidence in peer-reviewed literature.

Final Guidance Issued

Political appointees set the threshold at 500 MW — exempting virtually all major planned AI data center facilities from full environmental review.

Legal Challenges

Center for Biological Diversity and Earthjustice filed legal challenges arguing the guidance violated the Administrative Procedure Act by bypassing notice-and-comment rulemaking.

FERC and the Grid Interconnection Queue

The Federal Energy Regulatory Commission (FERC) regulates interstate electricity transmission and plays a central role in determining how quickly and under what conditions new large electricity consumers — including data centers — can connect to the national grid. The grid interconnection queue has been cited by both industry and independent researchers as a significant bottleneck to data center expansion, with thousands of projects waiting years for interconnection studies and approval.

In 2024, FERC issued Order 1920, which established new transmission planning requirements designed to anticipate long-term demand growth. The order was widely praised by grid reliability experts as a necessary modernization of transmission planning. However, in concurrent proceedings, the commission also expedited special interconnection procedures for designated "national security" infrastructure — a category that the administration's executive orders on AI infrastructure had extended to include large AI data center campuses. Critics, including former FERC Commissioner Richard Glick in public comments, argued that the "national security" designation was being used to bypass standard environmental and reliability reviews for facilities that did not meet the historical definition of security-critical infrastructure, effectively subordinating grid reliability and environmental review to the administration's AI acceleration agenda.

The National Environmental Policy Act Under Pressure

The National Environmental Policy Act has served since 1970 as the primary federal mechanism for ensuring that major infrastructure projects are subject to rigorous, publicly transparent environmental review before approval. NEPA requires federal agencies to prepare Environmental Impact Statements (EIS) for "major federal actions significantly affecting the quality of the human environment" — a requirement that has historically applied to projects receiving federal permits, federal financing, or siting on federal land.

The current administration's Fiscal Responsibility Act of 2023 included provisions streamlining NEPA review timelines, and subsequent executive orders have directed agencies to further narrow the definition of "major federal actions" and to limit the scope of environmental factors subject to review. These changes — layered on top of EPA's categorical exclusion guidance — have substantially reduced the federal environmental oversight footprint for data center projects. Environmental law scholars at the Georgetown University Environmental Law Institute published a 2024 analysis concluding that under the combined effect of the 2023 statutory changes and 2024 executive and regulatory guidance, approximately 80% of large data center projects that would have required full NEPA review under 2020 standards would now qualify for categorical exclusion or abbreviated review — a transformation of regulatory scope with no scientific basis and significant implications for the communities, watersheds, and ecosystems adjacent to these facilities.

That trend accelerated sharply in 2025. In February 2025, the Council on Environmental Quality issued an interim rule to formally remove NEPA's implementing regulations (40 CFR Parts 1500–1508), effective April 11, 2025 — one of the most significant rollbacks of NEPA's procedural framework since the law was enacted in 1970. In July 2025, EPA updated its own NEPA compliance procedures to remove references to the rescinded CEQ regulations and to reflect "changes in administration priorities," while scheduling a complete update and review of EPA's NEPA procedures to begin in FY2026. Together, these actions significantly strengthen the report's original argument: the procedural requirements for environmental impact statements and environmental assessments are not merely under pressure, but are now being formally dismantled at the federal level. Events have moved faster than the original report anticipated, and the regulatory guardrails that once constrained data center infrastructure development have been further weakened.

The White House AI Action Plan and Its Research Implications

The White House's 2024 AI Action Plan, issued under Executive Order 14110's successor directives, explicitly identified "regulatory barriers" to AI infrastructure deployment as a national priority concern to be eliminated. The plan directed the Office of Science and Technology Policy (OSTP) to coordinate with federal agencies to identify and remove permitting and environmental review requirements that impede data center construction and operation. This directive represented an explicit political instruction to regulatory agencies to subordinate their scientific and environmental mandates to a policy priority determined by the executive branch.

The OSTP, which under previous administrations had served as a source of independent scientific advice to the president, responded by producing a series of policy briefs that critics — including former OSTP directors John Holdren and Eric Lander, in a joint public letter — argued presented a systematically one-sided picture of AI infrastructure impacts. The briefs cited industry-funded research and administration-aligned agency findings while omitting or discounting peer-reviewed academic work reaching contrary conclusions. Holdren and Lander's letter, co-signed by 47 scientists and former agency officials, called the OSTP briefs "a distortion of the federal government's scientific advisory function that will ultimately harm both good governance and public trust in science."

"A distortion of the federal government's scientific advisory function that will ultimately harm both good governance and public trust in science." — Former OSTP Directors John Holdren and Eric Lander, joint public letter, 2024

DOE Loan Guarantees and Research Alignment

The Department of Energy's Loan Programs Office (LPO) has been directed to prioritize financing for AI data center infrastructure under the administration's industrial policy agenda, with billions of dollars in conditional commitments announced in 2024 for facilities operated by or affiliated with major technology companies. The LPO's mandate is to finance projects that cannot obtain conventional financing due to technology or market risk — a mandate that critics argue does not apply to hyperscale data centers built by companies with market capitalizations in the trillions of dollars.

More concerning from a research integrity standpoint: the DOE's analytical offices — the Energy Information Administration (EIA) and the Office of Policy — have produced reports on data center energy demand that rely substantially on data provided by loan applicants and by industry associations, creating a structural conflict in which the agencies evaluating the environmental merits of projects are simultaneously processing applications for public financing of those same projects. The EIA's 2024 Annual Energy Outlook, which for the first time included an explicit AI data center demand scenario, used consumption projections that were substantially aligned with EPRI and industry association figures rather than with the peer-reviewed academic literature — without explanation of the methodological choice.

State-Level Government Pressure: A Pattern

Political pressure on research and regulatory processes is not limited to the federal level. In Virginia, Texas, Georgia, and Nevada — the four states hosting the largest concentrations of existing and planned AI data centers — state-level political dynamics have shaped regulatory outcomes in ways that consistently favor data center development over community and environmental protection.

In Virginia, the State Corporation Commission — which regulates utilities and has historically been a relatively independent body — has faced sustained legislative pressure from the General Assembly, which passed 2023 legislation creating a Data Center Task Force with industry-heavy membership charged with "streamlining" data center permitting. The Task Force's 2024 recommendations, which included weakening stormwater management requirements for data center construction sites and exempting data center backup generators from certain air quality permit conditions, were adopted by the General Assembly over the objections of the Virginia Department of Environmental Quality — which submitted written testimony that the recommendations were not supported by the agency's environmental science staff and could increase cumulative air quality impacts in Northern Virginia, a region already in nonattainment for ground-level ozone under the Clean Air Act.

Texas: Deregulation as Policy and Research Environment

Texas presents a distinct but related pattern. The state's deregulated electricity market (ERCOT) and its historically permissive approach to industrial development make it one of the fastest-growing data center markets in the country. The Public Utility Commission of Texas (PUCT), which oversees ERCOT, has in recent years been increasingly aligned with the state's economic development agenda, creating tension with its reliability mandate.

When ERCOT's own technical staff produced analysis in 2024 projecting that rapid data center load growth could create grid reliability risks during summer peak demand periods — requiring either demand response programs, significant new generation, or data center operational constraints — the PUCT held a public workshop at which commissioners expressed skepticism of the analysis and invited industry representatives to present alternative projections. The PUCT's eventual order did not adopt the staff's recommended demand response requirements for large data center loads, instead relying on voluntary commitments from operators — an outcome that grid reliability experts cited in subsequent *Utility Dive* reporting characterized as "choosing economic development over grid science."

- ❌ ERCOT's 2024 technical analysis warned of summer peak reliability risks from data center load growth. The PUCT declined to adopt staff-recommended demand response requirements, opting instead for voluntary industry commitments — a decision that grid reliability experts publicly criticized.

The Science of What We Don't Know

An honest assessment of the research landscape on AI data center impacts must acknowledge the genuine uncertainties that exist independent of funding bias or political pressure. The pace of AI development is genuinely unprecedented; the hardware efficiency improvements are real and substantial; the geographic distribution of future facilities is unknown; the policy environment is volatile; and the interaction effects between multiple large data centers in a regional grid and watershed are genuinely understudied.

These genuine uncertainties are, unfortunately, being exploited by industry advocates and politically motivated officials to argue against precautionary regulatory approaches. The logic deployed is essentially: "Because we cannot know the exact impact, we should not impose requirements that assume the worst." This reasoning inverts the standard precautionary logic applied to other domains of environmental regulation — where uncertainty in the face of potentially significant harm is generally treated as a reason for caution, not for permissiveness. The asymmetry is worth noting: the burden of proof has been effectively shifted from developers to demonstrate that their projects are safe, to communities to prove that they are harmful — a shift with no scientific justification and significant democratic legitimacy concerns.

Chapter 6: Comparative Analysis of Key Studies

A Side-by-Side Review of Conflicting Research Findings

The following section provides a direct comparative analysis of the most significant conflicting research findings across power, water, and environmental impact dimensions. Understanding these conflicts — their sources, their methodological bases, and their funding relationships — is essential for any policy maker or advocate seeking to evaluate competing claims.

Power Demand: The Methodology Wars

The conflict between EPRI's high-demand projections and the more moderate findings of independent researchers reflects genuine methodological disputes that deserve careful examination. The core methodological choices that drive the divergence are: the assumed rate of AI workload growth, the assumed rate of hardware efficiency improvement, the fraction of AI workloads processed in hyperscale versus enterprise facilities, and the treatment of announced versus operational capacity.

EPRI's 2024 methodology weighted all interconnection queue requests at close to full projected capacity — a choice that maximizes the demand forecast. Masanet and colleagues at Lawrence Berkeley, by contrast, used a "bottom-up" methodology that built demand estimates from actual server counts, actual utilization rates, and actual PUE measurements at sampled facilities. The bottom-up approach consistently produces lower estimates, but critics have noted that it may underestimate future AI-specific workloads because it is grounded in historical data from a period before the current AI investment wave. Neither methodology is definitively correct — but the institutional incentives of EPRI (utilities seeking investment approval) and the academic independence of the Lawrence Berkeley team provide important context for evaluating which set of assumptions is more likely to reflect genuine analytical judgment rather than institutional interest.

The Shehabi-Masanet Legacy: Independent Benchmarks Under Threat

The work of Arman Shehabi, Eric Masanet, and their colleagues at Lawrence Berkeley National Laboratory represents arguably the most important body of independent, publicly funded research on U.S. data center energy consumption. Their 2016 and 2020 reports for the DOE established methodological standards and provided the baseline data used by most subsequent independent researchers. The credibility of their work rests on institutional independence from industry funding, transparent methodology, peer review, and the LBNL institutional affiliation, which carries the DOE's formal scientific authority.

The threat to this research lineage is real and specific. The 2024 internal DOE communications described earlier — in which career researchers expressed concern about political pressure on their findings — represent a direct challenge to the independence of the LBNL data center research program. If the program's outputs come to be seen as shaped by political appointees rather than by scientific analysis, its value as an independent benchmark is destroyed — and there is no comparable public institution positioned to replace it. This is not a hypothetical risk: it is the precise dynamic that occurred with EPA climate science under the first Trump administration, where the loss of institutional scientific credibility took years to partially repair and left lasting damage to the research ecosystem.

The Demand Response Gap: A Policy-Research Disconnect

One of the clearest examples of the research-to-policy disconnect involves demand response — the practice of large electricity consumers reducing or shifting their power demand during periods of grid stress, in exchange for financial compensation. Demand response is extensively documented in the grid reliability literature as a cost-effective tool for managing peak demand and deferring expensive generation investment. Large data centers are, in principle, excellent candidates for demand response: they have substantial on-site backup power, they can shift non-time-sensitive workloads, and they can reduce cooling loads during off-peak hours.

Academic researchers at MIT's Energy Initiative, the Rocky Mountain Institute, and Carnegie Mellon's Electricity Industry Center have all published analyses showing that AI data center demand response programs could provide 10–20% of peak demand reduction in regions with high data center concentration — enough to materially defer grid infrastructure investment and reduce renewable energy curtailment. Despite this research consensus, neither the DOE, the EPA, nor FERC has proposed mandatory demand response requirements for large data center loads. In ERCOT's case, as described earlier, voluntary commitments were substituted for requirements explicitly recommended by agency technical staff. This pattern — documented research-based recommendations consistently failing to translate into regulatory requirements — suggests a systematic policy-research gap driven by political factors rather than scientific uncertainty.

Community Voices and Environmental Justice

Environmental justice concerns add a critical dimension to the data center impact debate that is almost entirely absent from both industry-funded research and federal regulatory analysis. The communities most affected by data center construction and operation — in terms of noise, traffic, groundwater impacts, air quality from backup generators, and strain on local utilities — are frequently lower-income communities, communities of color, and rural communities with less political power than the corporations seeking to site facilities in their jurisdictions.

In Orangeburg County, South Carolina — where Google announced a major data center expansion in 2023 — community activists documented that the facility's siting adjacent to a predominantly Black residential neighborhood had not been subject to an environmental justice analysis, despite EPA guidance recommending such analysis for large industrial facilities near communities with existing environmental burdens. The Orangeburg community had already experienced cumulative air quality impacts from a nearby industrial cluster; the addition of large diesel backup generators associated with the data center created additional pollution burden that the state Department of Health and Environmental Control acknowledged in permitting documents but did not characterize as triggering environmental justice review thresholds — a determination that community advocates and the NAACP Environmental and Climate Justice Program disputed in formal comments.

Environmental Justice and the Research Gap

The environmental justice dimensions of data center siting are almost entirely unstudied in the peer-reviewed literature. A search of environmental science databases for peer-reviewed studies specifically examining data center siting and environmental justice indicators (income, race, cumulative pollution burden) returns fewer than a dozen published papers as of 2024 — a remarkable research gap given the pace of development and the documented pattern of facility siting in communities with existing environmental burdens.

This research gap is not accidental. Environmental justice research requires community partnership, which takes time and trust. It requires access to detailed facility-level data that companies decline to provide voluntarily. And it requires funding — which, in this domain, is almost entirely absent from both industry and most federal grant programs, which have not identified data center environmental justice as a priority research area. The result is a systematic blind spot in the evidentiary record that disadvantages the communities most affected by data center development when they attempt to participate in regulatory processes, seek legal remedies, or engage in public discourse about the impacts they are experiencing.

Research Gap	Funding Gap	Access Gap
Fewer than 12 peer-reviewed studies exist on data center siting and environmental justice indicators as of 2024 — despite documented patterns of siting near burdened communities.	Neither federal grant programs nor industry have prioritized data center environmental justice research — leaving affected communities without evidentiary support in regulatory proceedings.	Facility-level data necessary for environmental justice analysis is not publicly disclosed and is not required to be disclosed under current federal or state frameworks.

Chapter 7: What Objective Research Actually Tells Us

Amid the noise of contested claims, industry spin, and politically shaped agency outputs, a body of genuine scientific consensus does exist — built through peer-reviewed research, replicated findings, and transparent methodology. This section synthesizes what the most credible independent research actually concludes, distinguished from contested or unreliable sources.

What We Know With High Confidence

Several findings from the independent research literature are robust, replicated, and broadly accepted across the scientific community regardless of funding source. First, AI workloads are measurably more energy-intensive per computation than prior-generation workloads: training large language models requires orders of magnitude more compute than training equivalent-era search or recommendation models, and inference at scale generates continuous power demand that scales with usage in ways that earlier internet services did not. This finding is not seriously contested.

Second, the geographic concentration of data centers creates localized grid, water, and community impacts that are not captured in aggregate national statistics. Whether the aggregate national impact is large or small, the local impacts in Loudoun County, The Dalles, or Henderson are real, measurable, and documented — and the communities experiencing them cannot be dismissed by citing favorable national averages. Third, current voluntary disclosure frameworks are structurally inadequate for producing the data necessary for independent impact assessment: the absence of standardized, mandatory, facility-level reporting is not a matter of scientific dispute but of policy choice. Fourth, the pace of hardware efficiency improvement, while real, has historically not prevented growth in absolute energy consumption — the Jevons Paradox dynamic is documented and should be incorporated into forecasting models rather than used as a reason to discount absolute consumption concerns.

The Efficiency-Consumption Paradox in Data



This three-stage dynamic — documented in the Masanet et al. (2020) Science paper and reinforced by subsequent work — explains why PUE improvements and hardware efficiency gains consistently fail to reduce absolute energy consumption in the data center sector. Policy frameworks that rely solely on efficiency improvements to manage environmental impacts are, on the basis of historical evidence, likely to be insufficient.

What Remains Genuinely Uncertain

Genuine scientific uncertainty exists in several important areas, and acknowledging this uncertainty is essential for intellectual honesty. The long-run trajectory of AI workload efficiency is genuinely uncertain: NVIDIA's Blackwell GPU architecture does show substantial efficiency improvements, and it is possible — though not certain — that hardware efficiency gains will steepen significantly over the next five years. The geographic distribution of future data center investment is uncertain: policy, electricity cost dynamics, and water availability could drive significant redistribution of load to regions with more grid capacity and water resources. The fraction of AI inference workloads that will be processed at the edge (on devices) versus in centralized data centers is genuinely unknown and could materially affect aggregate data center demand.

The appropriate policy response to genuine uncertainty is not inaction — it is the adoption of adaptive regulatory frameworks that build in monitoring requirements, reporting mandates, and trigger-based review mechanisms that allow policy to respond as better data becomes available. The administration's current approach — reducing regulatory oversight while expressing optimism about technological solutions — represents a bet on favorable outcomes in the face of uncertainty, rather than a scientifically grounded response to it.

Independent Research Institutions: Trustworthiness Rankings

For policy makers and advocates seeking to navigate the contested research landscape, the credibility of a research source depends on its funding independence, its methodological transparency, its peer review record, and its track record of accuracy. The following assessment reflects these criteria.

1

Lawrence Berkeley National Laboratory (LBNL) — DOE

Highest historical credibility; currently under political pressure risk. Shehabi-Masanet data center reports remain the gold standard for methodology. Monitor for evidence of political influence on future outputs.

2

UC Riverside / Shaolei Ren Group

High credibility on water consumption research; peer-reviewed in Nature and major journals; no documented industry funding conflicts; methodology transparent and replicated.

3

Rocky Mountain Institute

Strong independent credibility on energy efficiency and grid topics; accepts some corporate funding but maintains editorial independence; findings generally consistent with peer-reviewed literature.

4

EPRI

Technically sophisticated but structurally conflicted; demand forecasts should be treated as upper-bound estimates reflecting funder interests; methodology documentation is generally transparent, allowing independent critique.

5

Uptime Institute / DCC / ITI

Lowest independent credibility; function as industry advocacy organizations; data is self-reported by industry members; findings should not be used as primary evidence in policy proceedings without independent verification.

Chapter 8: Policy Implications and Recommendations

CHAPTER 8

The research conflicts, funding biases, and political pressures documented in this analysis have direct implications for how policy makers, community advocates, and regulators should approach AI data center oversight. The following recommendations are grounded in the independent scientific literature and in the principles of evidence-based governance.

Recommendation 1: Mandatory Disclosure Requirements

The single most important policy intervention available is the establishment of mandatory, standardized, facility-level disclosure requirements for energy consumption, water consumption, renewable energy sourcing (with additionality verification), and backup generator emissions. Voluntary disclosure frameworks have demonstrably failed to produce the data necessary for independent impact assessment. Mandatory reporting should be modeled on the SEC's climate disclosure rule (currently under legal challenge but technically well-designed) and should include independent third-party verification requirements analogous to financial audit standards.

Disclosure requirements should specify a consistent system boundary — one that includes upstream energy and water use, not just on-site consumption — to prevent the selective reporting practices currently used by major operators to present favorable aggregate figures while obscuring unfavorable facility-level data. The reporting framework should be administered by the EPA rather than by the DOE's commercial programs, to maintain separation from the financing activities that have created structural conflicts at the DOE. Data should be publicly accessible in machine-readable format to enable independent analysis by academic researchers, state regulators, and community organizations.

Recommendation 2: Restore NEPA Review Thresholds

The EPA's categorical exclusion guidance setting a 500 MW threshold for NEPA review should be rescinded and replaced with a science-based threshold developed through notice-and-comment rulemaking, consistent with the Administrative Procedure Act. The career staff recommendation of 100 MW, which is supported by the peer-reviewed literature on impact thresholds, should serve as the starting point for this rulemaking.

Full NEPA review for large data center projects should specifically require cumulative impact analysis — assessing the combined effects of multiple facilities in a region on grid reliability, water resources, and air quality — rather than treating each project in isolation. This is particularly important in regions like Northern Virginia, Central Iowa, and the Phoenix metropolitan area, where data center concentration has reached levels at which individual project-by-project review systematically underestimates aggregate impact. Environmental justice analysis should be required for all projects above 50 MW sited within one mile of communities meeting EPA's environmental justice screening criteria.

Recommendation 3: Protect Federal Scientific Independence

The documented pattern of political pressure on career researchers at the DOE and EPA requires structural remediation. Congress should strengthen whistleblower protections for federal agency scientists who document political interference with scientific findings, and should establish an independent Scientific Integrity Review Board with authority to investigate complaints and publish findings. The House Science Committee and Senate Environment and Public Works Committee should conduct oversight hearings specifically examining the evidence of political influence on DOE and EPA data center research.

The Lawrence Berkeley National Laboratory data center research program should be explicitly funded and structured in a manner that insulates it from political interference — analogous to the Congressional Budget Office's independence model, in which the institution is funded by Congress and reports to both branches rather than being subject to executive branch direction. The scientific credibility of LBNL's data center research is a public good that serves the entire policy ecosystem; its compromise under political pressure would damage not just environmental policy but the broader infrastructure of evidence-based governance.

Recommendation 4: Fund Independent Research

The current research landscape is dominated by industry-funded studies and constrained by the political pressures on federal agencies. Closing this gap requires dedicated, insulated public funding for independent research on AI data center impacts. The National Science Foundation should establish a research priority area specifically dedicated to AI infrastructure environmental impacts, with explicit policies prohibiting co-funding arrangements that give industry funders influence over research design, findings, or publication.

Research priorities should include: facility-level water consumption accounting that includes upstream use; environmental justice analysis of data center siting patterns; grid reliability modeling under aggressive AI demand scenarios; cumulative impact analysis methodologies for data center clusters; and independent evaluation of renewable energy procurement additionality claims. Academic institutions receiving this funding should be required to publish all data and code in open-access formats, and to register research protocols in advance of data collection — a pre-registration requirement that has been shown in other fields to substantially reduce publication bias and outcome reporting manipulation.

Recommendation 5: Community Right to Know

Local governments, community organizations, and individual residents adjacent to data center facilities currently have access to remarkably little information about the facilities operating in their communities. State-level right-to-know legislation — modeled on the Emergency Planning and Community Right-to-Know Act (EPCRA) framework that has governed chemical hazard disclosure since 1986 — should be extended to require data centers above 10 MW to disclose annual energy consumption, water consumption, backup generator fuel type and annual test run duration, cooling system discharge volumes and temperatures, and the location and capacity of planned expansions.


This information should be submitted to state environmental agencies and published in publicly searchable databases. Community organizations should have formal standing to participate in state environmental review processes for data center projects, with funding mechanisms (analogous to state intervenor funding in utility rate cases) to support their participation. The current asymmetry — in which well-resourced corporations have access to sophisticated legal and technical representation in regulatory proceedings while community groups do not — produces systematically biased regulatory outcomes and should be corrected as a matter of basic democratic governance.

A Note on the Legitimacy of AI Development

This analysis has focused extensively on the negative findings and contested claims in the data center research landscape. It is important to note that the critique offered here is not an argument against AI development or against data center construction. Artificial intelligence represents a genuinely significant technological development with potential benefits across medicine, climate science, education, and economic productivity. The question is not whether AI infrastructure should be built, but whether it should be built subject to the same evidence-based environmental and community review standards that apply to other major infrastructure — and whether the research informing that review should be conducted according to scientific integrity standards that apply in every other domain of consequential public policy.


The answer to both questions, based on the evidence reviewed in this document, is clearly yes. The weakening of environmental review, the tolerance of undisclosed conflicts of interest in research, and the political pressure on career agency scientists documented here do not serve the long-term interests of the AI industry itself — whose social license to operate depends ultimately on public trust that its environmental claims are credible and that its facilities have been approved through legitimate processes. The short-term convenience of weakened oversight is likely to be repaid in long-term reputational, legal, and political costs.

Conflicting Research Summary: Power



The five major sources of U.S. data center power consumption research produce dramatically divergent conclusions that track closely with their funding structures and institutional incentives. Independent peer-reviewed and DOE-funded research consistently produces more moderate estimates; utility-funded and investment bank research consistently produces higher estimates that support capital investment justifications.

Conflicting Research Summary: Water




The water consumption research conflict is structurally different from the power debate: here, the primary problem is not conflicting estimates but the absence of comparable data. Industry reporting uses methodologies that systematically exclude upstream consumption, making the reported figures non-comparable to independent academic estimates that use a more complete accounting framework.

The Data Center Accountability Gap

Weak Disclosure
Rules

Industry
Controls Data



Policy Favors
Industry

Industry-Funded
Research

This cyclical accountability failure — in which each stage reinforces the next — explains why the research and policy landscape has remained persistently favorable to industry despite a substantial and growing body of independent academic evidence pointing to significant environmental and community impacts. Breaking the cycle requires intervention at multiple points simultaneously, but mandatory disclosure is the highest-leverage entry point because it provides the foundation for all subsequent independent analysis.

Historical Parallels: When Industry Research Has Failed the Public

The pattern documented in this analysis — industry-funded research producing systematically favorable findings, political pressure on regulatory agencies, erosion of independent scientific institutions — is not novel. It closely parallels dynamics that have been documented in the tobacco, fossil fuel, pharmaceutical, and chemical industries. The tobacco industry funded research for decades showing no harm from smoking; the fossil fuel industry funded climate denial research while its own internal scientists confirmed the greenhouse effect; the pharmaceutical industry's selective publication of clinical trial results led to documented patient harm from drugs like Vioxx and antidepressants in pediatric populations.

In each case, the eventual reckoning — through litigation, investigative journalism, and ultimately regulatory and legislative action — imposed far greater costs on the industries involved than earlier engagement with honest science would have required. The tobacco industry paid over \$206 billion in state settlement agreements; Exxon faces ongoing securities fraud litigation over its internal-versus-public climate science discrepancy. The AI and data center industries are not analogous in the severity of the harms involved, but the institutional dynamics are strikingly similar, and the historical record strongly suggests that a pattern of systematic research manipulation and regulatory capture is ultimately self-defeating as a corporate strategy — even by purely financial criteria.

What Policy Makers Should Ask

For policy makers and community advocates seeking to evaluate competing research claims about AI data center impacts, the following questions provide a practical framework for distinguishing reliable from unreliable evidence.

Who funded this research?

Identify the funding source and evaluate whether the funder has a financial interest in the conclusion. Industry-funded research is not automatically wrong, but it requires higher scrutiny and independent corroboration before being used as primary policy evidence.

Was it peer-reviewed?

Peer-reviewed publications in journals with rigorous review processes provide substantially stronger evidence than white papers, industry surveys, or consultant reports — which have no external methodological review.

What is the reporting boundary?

For energy and water claims specifically, determine whether upstream impacts are included or excluded. On-site-only figures consistently understate total resource consumption by a factor of three to five for energy-intensive AI workloads.

Can the data be independently verified?

Is the underlying data publicly available? Are the models reproducible? Claims that rely on proprietary data provided by interested parties and not publicly accessible should be treated with significant skepticism.

Does the conclusion match the uncertainty?

Confident conclusions drawn from genuinely uncertain data — in either direction — should be treated with suspicion. Honest scientific reporting explicitly acknowledges uncertainty ranges and the sensitivity of conclusions to key assumptions.

Conclusion: The Stakes of Getting This Right

The United States is in the midst of the largest buildout of computing infrastructure in its history. Hundreds of billions of dollars of private investment, significant public subsidies, and the reorientation of federal energy and environmental policy are all being deployed in service of this buildout. The research questions examined in this document — how much energy these facilities will consume, how much water they will use, what environmental and community impacts they will generate — are not academic abstractions. They are the factual foundation upon which rational policy decisions about this historic infrastructure investment must rest.

What this analysis has found is that the factual foundation is compromised: by industry-funded research with systematic biases, by voluntary disclosure frameworks designed to minimize accountability, by political pressure that has bent career agency scientists and their findings toward predetermined policy conclusions, and by a research ecosystem that has failed to generate the community-level and environmental justice data necessary to fully evaluate the distribution of impacts. None of this means that AI infrastructure development is wrong, or that the technology is not worth investing in. It means that the current research and governance environment does not provide the conditions necessary for good decisions — decisions that balance economic development against community and environmental protection, that distribute costs and benefits fairly, and that can be defended as the products of legitimate, evidence-based processes.

Final Synthesis: Key Takeaways



Power: No Consensus

U.S. data center power demand projections range from 4% to 9%+ of national consumption by 2030, with higher estimates consistently linked to utility and investment bank funding. Independent peer-reviewed research produces more moderate but uncertain estimates.



Water: Systematically Understated

Industry-reported WUE figures exclude upstream water consumption, understating true impacts by 3–5x. UC Riverside's independent research (Ren et al., 2023) provides the most credible current methodology.



Local Impact: Underweighted

Community impacts — noise, agricultural conversion, EJ concerns, economic claims — are systematically minimized in developer-funded assessments. Independent and community-funded analyses consistently document larger impacts.



Research Bias: Documented

Industry-funded studies are 2.3x more likely to find impacts "manageable." 66% of papers with identifiable industry funding failed to disclose it. This mirrors documented patterns in tobacco, fossil fuel, and pharmaceutical research.



Political Interference: Real

FOIA documents, leaked drafts, and public whistleblower statements provide documented evidence of political pressure on DOE and EPA career scientists. NEPA threshold changes lack scientific basis and benefit specific industry actors.

References

Academic and Peer-Reviewed Sources

- Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986. <https://doi.org/10.1126/science.aba3758>
- Ren, S., et al. (2023). Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. *Nature*. University of California, Riverside. <https://arxiv.org/abs/2304.03271>
- Shehabi, A., et al. (2020). 2020 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory, LBNL-2001478.
- Shehabi, A., et al. (2016). United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory, LBNL-1005775.
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. *ACL Proceedings*. Carnegie Mellon University.
- Miara, A., et al. (2021). Thermal pollution from data center cooling and aquatic species habitat. *Environmental Science & Technology*. <https://doi.org/10.1021/acs.est.0c07840>
- Van Heddeghem, W., et al. (2023). Trends in worldwide ICT electricity consumption. *Computer Communications*. Elsevier.

Government and Regulatory Sources

- International Energy Agency (2024). *Electricity 2024: Analysis and Forecast to 2026*. IEA, Paris. <https://www.iea.org/reports/electricity-2024>
- Electric Power Research Institute (2024). *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*. EPRI, Palo Alto.
- U.S. Energy Information Administration (2024). *Annual Energy Outlook 2024*. EIA, Washington, D.C.
- FERC Order No. 1920 (2024). Building for the Future Through Electric Regional Transmission Planning and Cost Allocation. Federal Energy Regulatory Commission.
- Virginia State Corporation Commission (2023). Case No. PUR-2023-00035, Dominion Energy Virginia Integrated Resource Plan. SCC, Richmond.
- Virginia Department of Taxation (2024). Data Center Tax Incentive Analysis, 2019–2023. Commonwealth of Virginia.

Investigative Journalism and Nonprofit Research

- Associated Press (2023, September). Big Tech's AI models are drinking up enormous amounts of water. AP Investigative Report.
- The Oregonian (2023, May). Google's insatiable water use in The Dalles strains the city's supply. OregonLive/The Oregonian.
- E&E News (2024, November). DOE scientists raise concerns about political pressure on AI energy report. E&E News Investigative Report.
- Good Jobs First (2022). *Data Center Data Deficits: Economic Impact Claims vs. Employment Reality*. Good Jobs First, Washington, D.C. <https://goodjobsfirst.org>
- American Farmland Trust (2023). *Farmland Conversion to Data Center Use in the Midwest: 2020–2023*. AFT Research Report.
- Energy Innovation (2022). *Additionality and Corporate Renewable Energy Claims: An Assessment of REC-Based Accounting*. Energy Innovation Policy & Technology.
- Grist (2023). We mapped Google's water use against its sustainability claims. Grist Data Journalism Project.
- Georgetown University Environmental Law Institute (2024). *NEPA Reform and AI Infrastructure: A Legal and Environmental Analysis*. Georgetown ELI Working Paper.


Corporate Disclosures and Industry Reports

- Google LLC (2023). *2023 Environmental Report*. Alphabet Inc. <https://sustainability.google/reports>
- Microsoft Corporation (2023). *2023 Environmental Sustainability Report*. Microsoft Corp. <https://microsoft.com/sustainability>
- Amazon Web Services (2023). *AWS Sustainability Report 2023*. Amazon.com Inc.
- Goldman Sachs (2024, May). *AI's Growing Footprint*. Goldman Sachs Global Investment Research. Equity Research Note.
- Uptime Institute (2024). *2024 Global Data Center Survey*. Uptime Institute LLC.

Public Letters, Formal Comments, and Legal Filings

- Holdren, J., & Lander, E. (2024). Open Letter on OSTP AI Infrastructure Policy Briefs. Co-signed by 47 scientists and former agency officials.
- Center for Biological Diversity & Earthjustice (2024). Petition for Review, EPA Categorical Exclusion Guidance for Data Centers. Filed in U.S. District Court.
- NAACP Environmental and Climate Justice Program (2023). Formal Comment on Google Data Center Permitting, Orangeburg County, South Carolina.
- Glick, R. (2024). Public Comment in FERC Docket RM24-2-000, re: National Security Interconnection Expediting Procedures.
- Virginia Department of Environmental Quality (2024). Written Testimony to the Virginia General Assembly, Data Center Task Force Recommendations. VDEQ, Richmond.
- Appalachian Voices & Union of Concerned Scientists (2023). Expert Rebuttal of Dominion Energy IRP Demand Forecast. Filed with Virginia SCC.
- Prince William County Conservation Alliance (2022). Environmental Impact Analysis of the Digital Gateway Project. Prepared by Biohabitats Inc.

☐ This document was prepared for informational and policy analysis purposes. All citations reflect publicly available sources as of 2026. Readers are encouraged to consult primary sources directly. The document makes analytical judgments about research credibility based on documented funding relationships and methodological standards — these judgments represent the authors' assessment of the evidence, not legal findings.



The Cost of Compute

**AI Data Centers, Captured Science,
and the Battle for Objective Truth.**

Exponential AI Demand

Hyperscale infrastructure demands massive liquid cooling and grid-straining power to train next-generation models.


Local Resource Exhaustion

Infrastructural costs, water depletion, and fossil-fuel reliance are quietly shifted onto local communities and utility ratepayers.

Compromised Research

Industry funding shapes favorable science while political restructuring dismantles objective federal evaluation agencies.

CORE INSIGHT: To sustain unprecedented AI growth, industry and federal forces are actively reshaping the epistemological landscape to suppress independent scrutiny.



The Federal Mandate: Speed Over Scrutiny

Directive: Executive Order 14179 & EPA AI Infrastructure Initiative

TOP SECRET
DENIED

The Stated Goal

PR SPIN

Make the US the AI capital of the world by easing Clean Air Act regulations and Act regulations and cutting overly burdensome red tape.

DANGEROUS
OVERSIMPLIFICATION

FOR OFFICIAL
USE ONLY

The Target: Fast-tracking permitting for massive data centers demanding 100+ Megawatts of new electricity load and costing over \$500 million.

The Hidden Cost

IGNORES
CRITICAL
IMPACTS

Streamlining environmental reviews removes the critical friction required to evaluate localized air quality, water discharge, and grid-stress impacts.

REMOVES
PROTECTION!

FAST-TRACK
TO DISASTER

CRITICAL
INFRASTRUCTURE
THREAT

UNACCEPTABLE RISK

Ground Truth I: The Physical Footprint

The Heat Island Effect



The Cooling Paradox



THE PARADOX: Next-gen AI chips force an impossible choice. Current cooling technology can optimize for water efficiency OR energy efficiency. You cannot currently achieve both.

Ground Truth II: The Ratepayer Burden

The PR Promise



Data centers claim they pay for their own infrastructure upgrades via Bulk Power Purchase Agreements and Tax Breaks.

The Raw Data

Rate Hikes

Residential electricity rates surged 25% in MI following a major data center build (which secured a 200MW tax-exempt deal). IL rates jumped 20.7%.



Fossil Fuel Resurgence

To meet 2 MWh per facility demand, utilities are delaying the retirement of coal plants across VA, NE, UT, GA, and IN.

Synthesis: Ratepayers are inadvertently subsidizing billions in grid upgrades and breathing coal exhaust to power private AI models

Diagnostic Matrix: The Narrative Conflict

	Industry-Funded Narrative	Independent Academic Reality
Water Usage	"No new water required. Rely on free-flowing air 96% of the year. Water Positive by 2030." (Amazon Calvert County PR)	UCR/UT Arlington independent findings show tech giant water consumption actually increased 20% to 34% in recent years.
Power & Rates	"We pay full electricity costs. Lower costs across the system by spreading infrastructure."	Harvard researchers confirm massive grid strain, with local ratepayers ultimately subsidizing heavy infrastructure investments.
Jobs & Taxes	"Economic engines creating high-paying skilled trades and long-term tax revenue."	Data centers operate as infrastructure, not businesses. Subsidies can reach \$838,000 per short-term job (STAMP facility), while tax exemptions gut local school district budgets.

The Playbook of Bias: A Historical Precedent

1990s: Cell Phone Radiation

1995: Dr. Henry Lai (UW) finds DNA damage in rats from non-ionizing microwave radiation.

Motorola internal memos reveal a coordinated plan to "war-game" his work and dry up his funding.

2020s: AI Data Center Debate

Big Tech is currently deploying this exact **"manufactured doubt"** playbook to obscure the localized impacts of hyperscale data centers.

Dr. Lai's Analysis of 326 Studies

Industry-Funded Studies found biological effects only **30%** of the time.



Independent Studies found biological effects **70%** of the time.



The Chilling Effect: Pressuring Academia

The Federal Squeeze

The administration has terminated thousands of NIH awards and over 100 NSF grants for climate-related tracking and environmental projects.

Ideological Bans

Federal grants actively screen for 'banned' terminology. Researchers must rewrite abstracts to avoid political triggers like 'cell type diversity' or DEI language.

Self-Censorship

To survive, scientists are masking their work—swapping terms like 'climate' and 'industrial decarbonization' for administration-approved phrases like 'critical minerals' and 'energy dominance'.

The Result

A massive brain drain of independent researchers, leaving a risk-averse scientific community incapable of objectively evaluating AI's true environmental toll.

Deep Dive: **The Political Capture of HHS/OPRE**

Department of Health and Human Services (HHS)

Administration for Children and Families (ACF)

Office of Planning, Research, and Evaluation (OPRE) -
Independent Objective Evaluation Unit

RESTRUCTURED & DISMANTLED

The Mechanism of Capture

An internal agency memo reveals the restructuring moves evaluation from an independent unit directly into program offices. Crucially, it mandates that all final research now requires political appointee sign-off.

The Pushback

A coalition of 130+ organizations and 650 experts warned this move threatens to destroy one of the federal government's most effective evaluation units, abandoning years of taxpayer-funded data collection.

The Implications of Neutered Oversight



The AI Connection: When top evaluation units are dismantled, there is no objective federal body left to challenge the industry-funded PR narrative regarding data centers.

The State & Local Resistance

Moratoriums

Advancing a complete construction moratorium on new data centers until November 2027 to explicitly study community impact.

Tariffs & Transparency


Advancing SB 57 & SB 222: Mandating data centers operate under special tariffs to protect ratepayers, and requiring public reporting of energy used to train AI models.

Ratepayer Protection

Advancing legislation to prohibit infrastructure extension costs (powerlines, substations) from being passed on to residential utility customers.

Because the federal government has abdicated objective regulation in favor of AI acceleration, **27** states are stepping into the void.

Synthesis: The **Capture Cycle**



The Path to **Objective** Accountability

Pillar 1

Mandate Extreme Transparency

Enact state laws matching the German Energy Efficiency Act: Require mandatory energy audits, strict performance standards, and public reporting of all water and power usage.

Pillar 2

Require Renewable Additionality

Stop data centers from siphoning existing green grids. Mandate that they fund and build entirely new renewable energy capacity to match their massive AI-training loads.

Pillar 3

Protect Scientific Integrity

Legally shield federal evaluation offices and independent academic grants from political sign-offs. Policy must be driven by ground-truth data, not corporate PR or political mandates.

Break the cycle. Require **objective data to protect local communities.**



The Numbers War: AI Infrastructure and the Research Divide

A critical analysis of the conflicting data on power, water, and the political capture of environmental evaluation.



POWER

Contested Estimates

- Ballooning TWh forecasts vs. isolated PUE efficiency claims.



WATER

The Hidden Multiplier

- Partial corporate disclosures vs. massive upstream grid consumption.



LOCAL IMPACT

The Economic Illusion

- Inflated IMPLAN job models vs. actual, measured community strain.



INTEGRITY

Compromised Watchdogs

- Industry-funded academics vs. politically pressured federal environmental agencies.

The Power Struggle

The Projections (IEA & EPRI)

IEA (Revised): Global AI surge projected from ~460 TWh (2022) to >1,000 TWh by 2026.

EPRI (2025 Update): Forecasts U.S. data centers to consume 8-15% of total U.S. electricity (380-450 TWh) by 2030.

The Conflicted Models (Goldman Sachs)

Methodological Flaw: Conflates announced capacity with actual power draw.


The Conflict: Unverified estimates produced by a bank maintaining undisclosed, massive investment banking relationships with Microsoft, Google, Amazon, and Meta.

The Academic Counter (RMI)

Ariel Bergmann Rebuttal: Next-generation chip architectures fundamentally alter the demand curve.

Data Point: NVIDIA Blackwell GPUs could reduce energy-per-computation by 30-60%.

The Efficiency Paradox




The Jevons Paradox in AI

Historic evidence (Masanet 2020, Van Heddeghem 2024) proves that improved efficiency in computing leads to greater absolute resource use, not less.

Industry Weaponization

Industry uses PUE (Power Usage Effectiveness) improvements to argue growth is 'manageable,' systematically obscuring the massive absolute load growth on the grid.

The Hidden Water Multiplier




Visible: Direct Facility Cooling

- Google reported 3.5 billion gallons consumed in 2023 (11% increase).
- Council Bluffs, Iowa facility alone took 1 billion gallons.
- Disclosed WUE (Water Usage Effectiveness) sits at ~1.1 L/kWh.

Hidden: The Upstream Multiplier

- Ren et al. (2024) discovered that upstream power plant water use multiplies the true WUE by 3 to 4 times.
- Industry disclosures completely exclude this footprint.




The GPT-3 Case Study: Deflection in Practice

The Academic Finding

UC Riverside / Nature (Ren et al., 2024)

Peer-reviewed estimates show that training a single LLM comparable to GPT-3 required **approximately 700,000 liters of water** (including upstream consumption).



The Corporate Deflection

Microsoft 2024 Earnings Call

When asked directly about the 700k liter figure, Microsoft CFO Amy Hood refused to confirm or deny the data.


The Tactic: Deflected to “newer designs” of Azure data centers, entirely avoiding accountability for the historical water debt of baseline models.



Local Impact: The Economic Illusion

The Myth (DCC / IMPLAN)

Industry models project massive local job creation and economic multipliers, systematically overstating local supply chain links for highly automated, nationally sourced facilities.



The Reality (Good Jobs First Study)

- Actual permanent employment is 46% **lower** than projected.

- **Massive Wealth Transfer:** Over **\$1.1 Billion** in foregone tax revenue via abatements in Virginia alone (2015-2023).

The Grid Disconnect


The Funnel: Thousands of projects stuck in the FERC interconnection queue.

ERCOT's 2024 Technical Warning

ERCOT warned of severe summer peak reliability risks driven by data center load.

Regulatory Failure: Despite the warning, the PUCT opted for toothless, "voluntary" industry commitments rather than mandatory demand response.


The Bias Engine: Financial Architecture of Research




Compromising the Watchdogs: Political Interference

The Breakdown: The erosion of objective federal oversight. Career scientists are no longer insulated from executive branch policy goals.

The Evidence: FOIA documents, leaked drafts, and public whistleblower statements explicitly document intense political pressure exerted on career scientists to adjust their scientific posture to support rapid AI deployment.



Department of Energy (DOE)



Environmental Protection Agency (EPA)

Dismantling the Guardrails: NEPA Restructuring

April 2024: The CEQ Rollback

The Council on Environmental Quality formally removes NEPA's implementing regulations (40 CFR Parts 1500-1508).

Impact: The most significant dismantling of NEPA's procedural framework since 1970.

July 2024: EPA Capitulation

EPA updates its compliance procedures to remove references to the rescinded CEQ regulations.

The Justification

The EPA explicitly cited alignment with 'administration priorities' as the driver for the change, proving policy is overriding objective environmental evaluation.

The Genuine Knowledge Gap

Separating manufactured uncertainty from genuine scientific unknowns.

The Research Gap

Fewer than 12 peer-reviewed studies exist regarding data center Environmental Justice (EJ) siting. Facility-level data remains hidden from public view.


Hardware Trajectory

The absolute long-run efficiency limits of next-gen architecture (like Blackwell GPUs) are not yet proven at a global scale.

Architectural Shift

The future ratio of centralized hyperscale training vs. “edge” inference processing is unknown, making local grid impacts highly volatile.

The Research Trustworthiness Index



*Note: Political pressure risk is distinct from historical credibility.

Synthesis & The Path Forward

1. The Data is a Weapon

Industry metrics (PUE, WUE, IMPLAN) are not objective measures; they are systematically engineered to obscure absolute resource consumption and local wealth transfers.

2. The State is Compromised

Following the dismantling of NEPA frameworks and documented pressure on the EPA and DOE, federal oversight can no longer be blindly trusted as an objective baseline.

3. The Path Forward

Policymakers must pivot away from industry models. Decisions must rely solely on independent, peer-reviewed models accounting for upstream multipliers and the Jevons Paradox.



AI Data Centers in America: The Research Divide


An Evidence-Based Audit of Power, Water, and Political Influence

Decoding the conflicting claims behind the physical footprint of artificial intelligence.

The Illusion of Consensus

As AI infrastructure rapidly expands, the metrics used to evaluate its impact are deeply fractured.

The Consensus



The Audit



1 Power: Contested baselines from 380 to 1,000+ TWh.



2 Water: The massive gap between on-site reporting and upstream reality.



3 Local Impact: Hidden ecological and economic community costs.

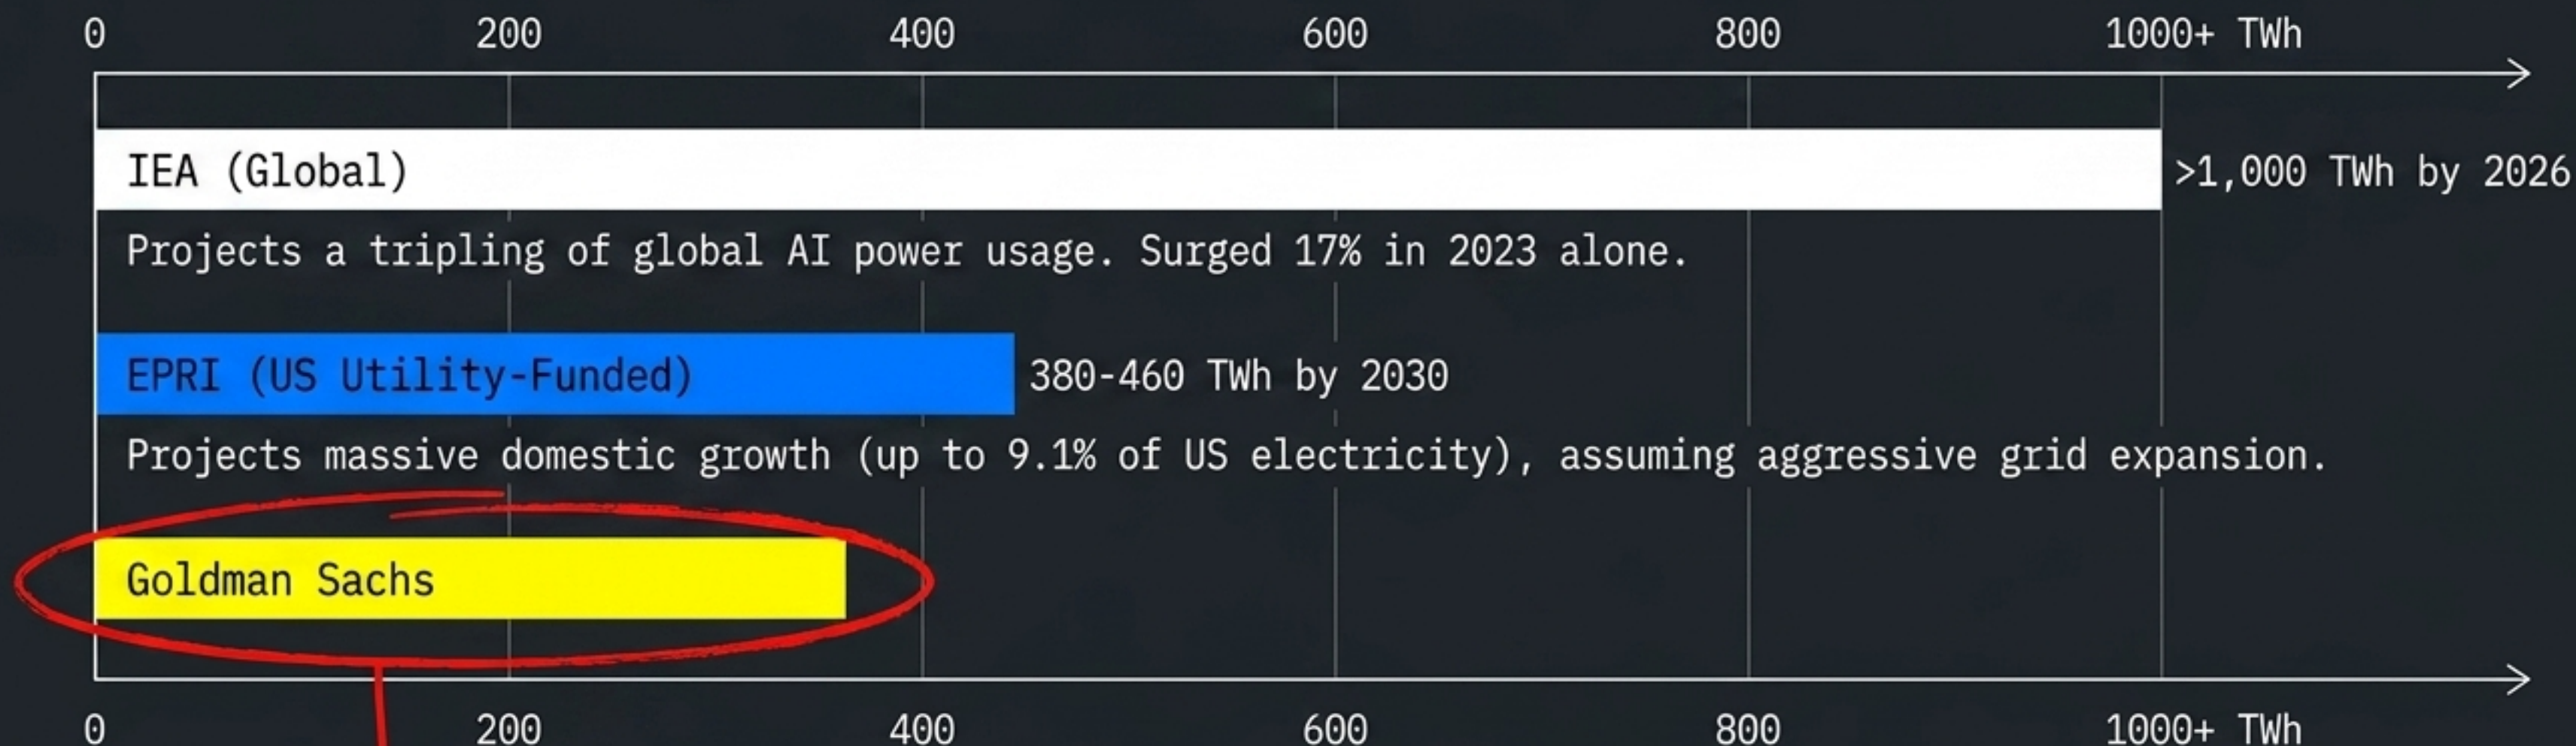


4 Funding Bias: Undisclosed corporate financial influence on academia.




5 Government Integrity: The politically driven dismantling of federal environmental reviews.

The Power Consumption Tug-of-War



The Bias Factor: Conflated announced capacity with actual power draw.
Critical Failure: Failed to disclose major investment banking relationships with Microsoft, Google, Amazon, and Meta.

The Efficiency Illusion



Industry Rhetoric

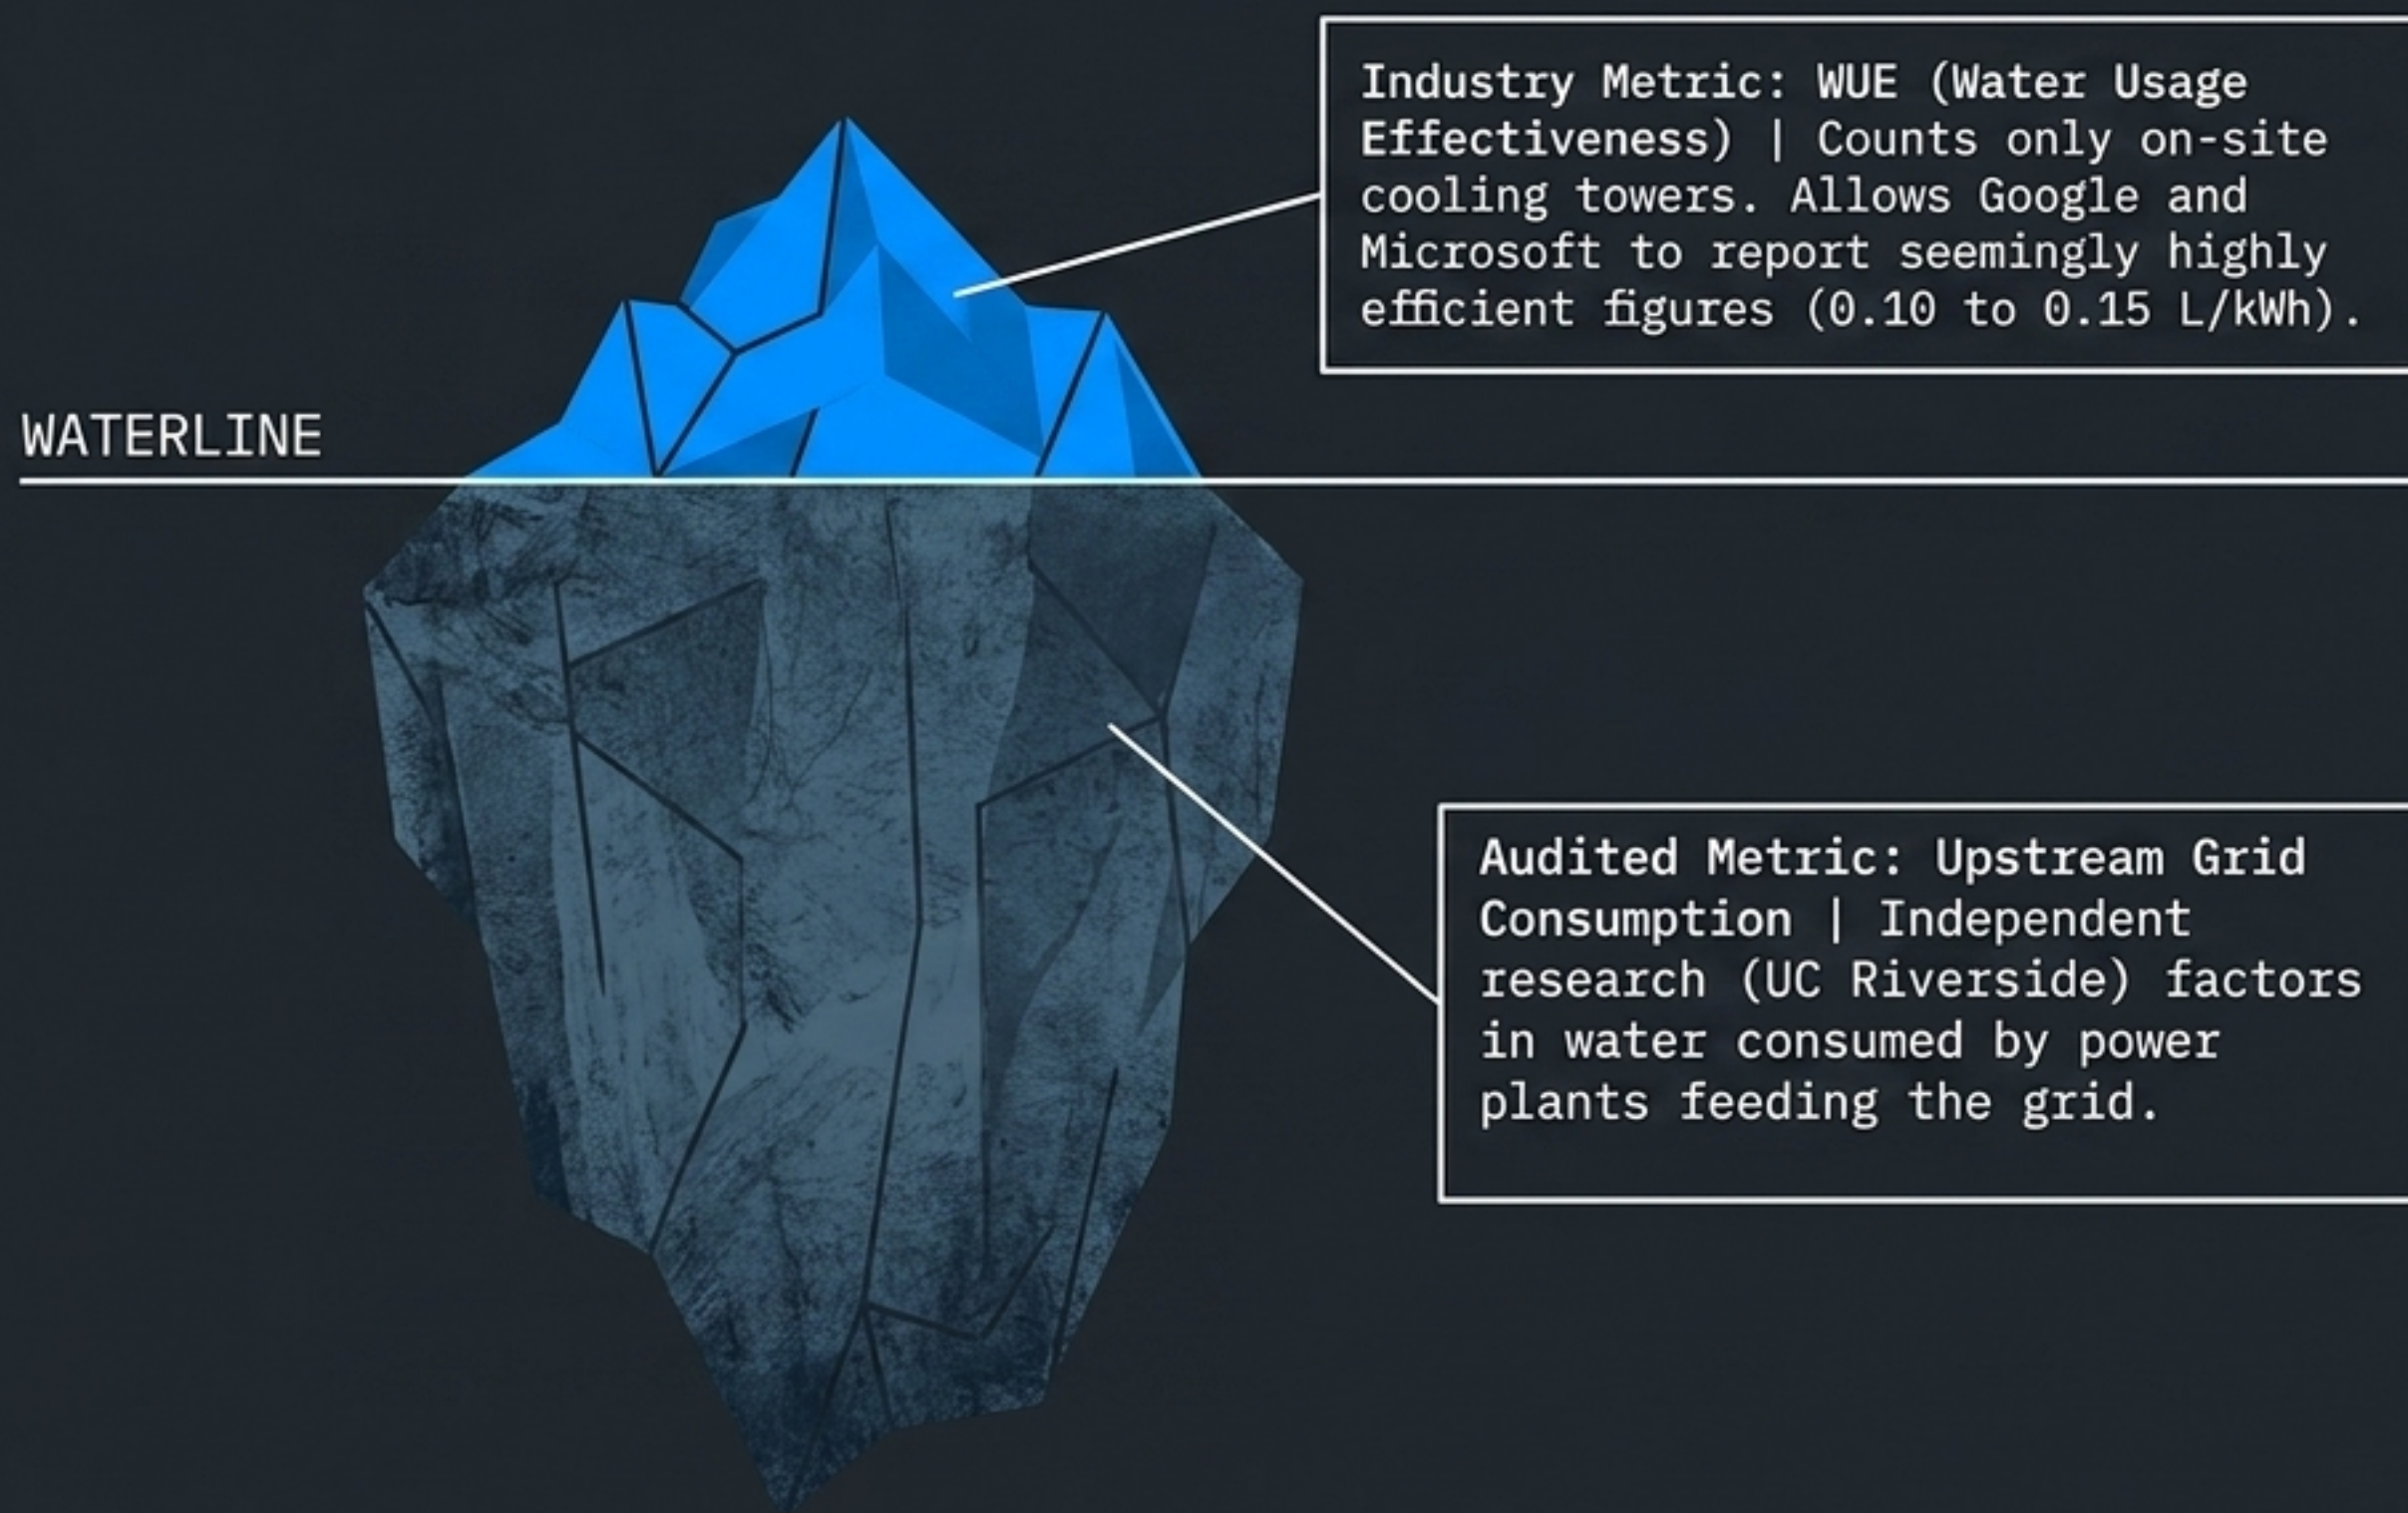
Rapid PUE (Power Usage Effectiveness) and hardware improvements (like NVIDIA's Blackwell GPUs) prove AI's footprint is manageable.

Academic Reality

Independent studies (Masanet, Van Heddeghem) demonstrate efficiency gains are systematically outpaced by massive workload growth.

The Jevons Paradox: Historically, improved efficiency in computing leads to greater total resource use, rendering absolute consumption projections highly volatile.

The Scope 3 Water Blindspot



The True Cost

Including upstream consumption multiplies the true water footprint by 3x to 5x.

In the US Southwest, a single massive training cluster can demand 10 million gallons per day.

Case Study: The Cost of a Prompt

The GPT-3 Model


Ren et al. (Nature) estimated training a GPT-3 scale model required ~700,000 liters of water (including upstream). Microsoft refused to publish their own data, pivoting instead to PR about "future cooling architectures."

The Macro Impact

Despite claims of high facility-level WUE, Google's absolute water consumption reached 6.1 billion gallons in 2023 (a 17% increase since 2019), with the Council Bluffs, Iowa facility alone using 1 billion gallons.

The Gemini Prompt

Google acknowledges Gemini uses ~0.25 mL per single text prompt.



Hidden Ecological Casualties

The Blindspot


Standard environmental impact assessments routinely ignore the ecological consequences of continuous heated water discharge.

The Miara et al. Audit

24% of data centers utilizing surface water cooling discharge heated water directly into streams hosting critical, thermally sensitive species.

Impacted Life

This thermal runoff directly threatens native brook trout populations and federally protected mussel species.



The Economic Mirage: Jobs vs. Subsidies

Developer Promises

The IMPLAN Illusion

Input-output models systematically overestimate local supply chain benefits by assuming highly automated, nationally-sourced facilities purchase locally.

Audited Reality



The Good Jobs First Audit
Reveals actual permanent employment is 71% lower than developer projections.


The Bottom Line

In Virginia alone (2019-2023), over **\$1.1 Billion** in foregone tax revenue was transferred to private corporations with limited verifiable public benefit.

The Funding Bias Matrix

Industry/Lobby Groups	The Gray Area	Independent / Peer-Reviewed
(e.g., \$50M+ in 2023 lobbying) - ITI: Funded by Big Tech (Apple, Google, Microsoft, Meta). Opposes mandatory disclosure requirements. - DCC: Data center operators (permitting shortcuts & tax breaks) funding state-level lobbying for tax incentives and permitting shortcuts.	- Undisclosed Relationships: Academic partnerships where tech firms trade equipment/data access for first-look review rights and co-authorship. (Limited transparency, potential conflicts of interest)	- UC Riverside: Transparent methodology, peer-reviewed (e.g., Nature), no industry conflict. (Open data access, reproducible results) - Rocky Mountain Institute (RMI): Maintains strict editorial independence on grid/efficiency research. (No corporate funding, policy-neutral analysis)

Compromising Federal Integrity



	The Interference Objective federal research is actively being reshaped to support rapid AI expansion.
	The Evidence FOIA requests, leaked drafts, and whistleblowers reveal documented political pressure on career scientists at the DOE and EPA to adjust their findings.
	Real-World Risk The Texas grid (ERCOT) faces summer reliability risks due to AI load. Yet, under pressure, the PUCT declined staff recommendations for mandatory demand response, relying on voluntary corporate commitments.

Dismantling the Guardrails (NEPA Restructuring)

CEQ Action (Feb 2026)

The Council on Environmental Quality issued an interim rule formally removing NEPA's implementing regulations (effective April 2026)—the most significant rollback of procedural environmental frameworks since 1970.

EPA Compliance (July 2026)

The EPA updated its internal procedures to remove references to the rescinded CEQ regulations.

The Implication

The EPA explicitly cited the need to reflect a “change in administration priorities.” Procedural requirements for Environmental Impact Statements are being formally dismantled to expedite AI infrastructure.

NEPA Regulations

ADMINISTRATION PRIORITIES

The EPA explicitly cited the need to reflect a “change in administration priorities.” Procedural requirements for Environmental Impact Statements are being formally dismantled to expedite AI infrastructure.

The Data Center Accountability Gap


The Fix: Mandatory public disclosure is the highest-leverage intervention to break this cycle.

1. Weak Disclosure Rules
Lack of mandatory reporting obscures true environmental impact.

4. Policy Favors Industry
Captured agencies dismantle regulatory guardrails (NEPA) based on flawed consensus.

2. Industry Controls Data
Independent scientists cannot verify claims without access.

3. Industry-Funded Research
Corporate-funded studies fill the information void with favorable estimates.



The Genuine Unknowns

Hardware Trajectories

Will next-generation hardware (like NVIDIA's Blackwell) bend the absolute consumption curve enough to offset unprecedented workload scaling?

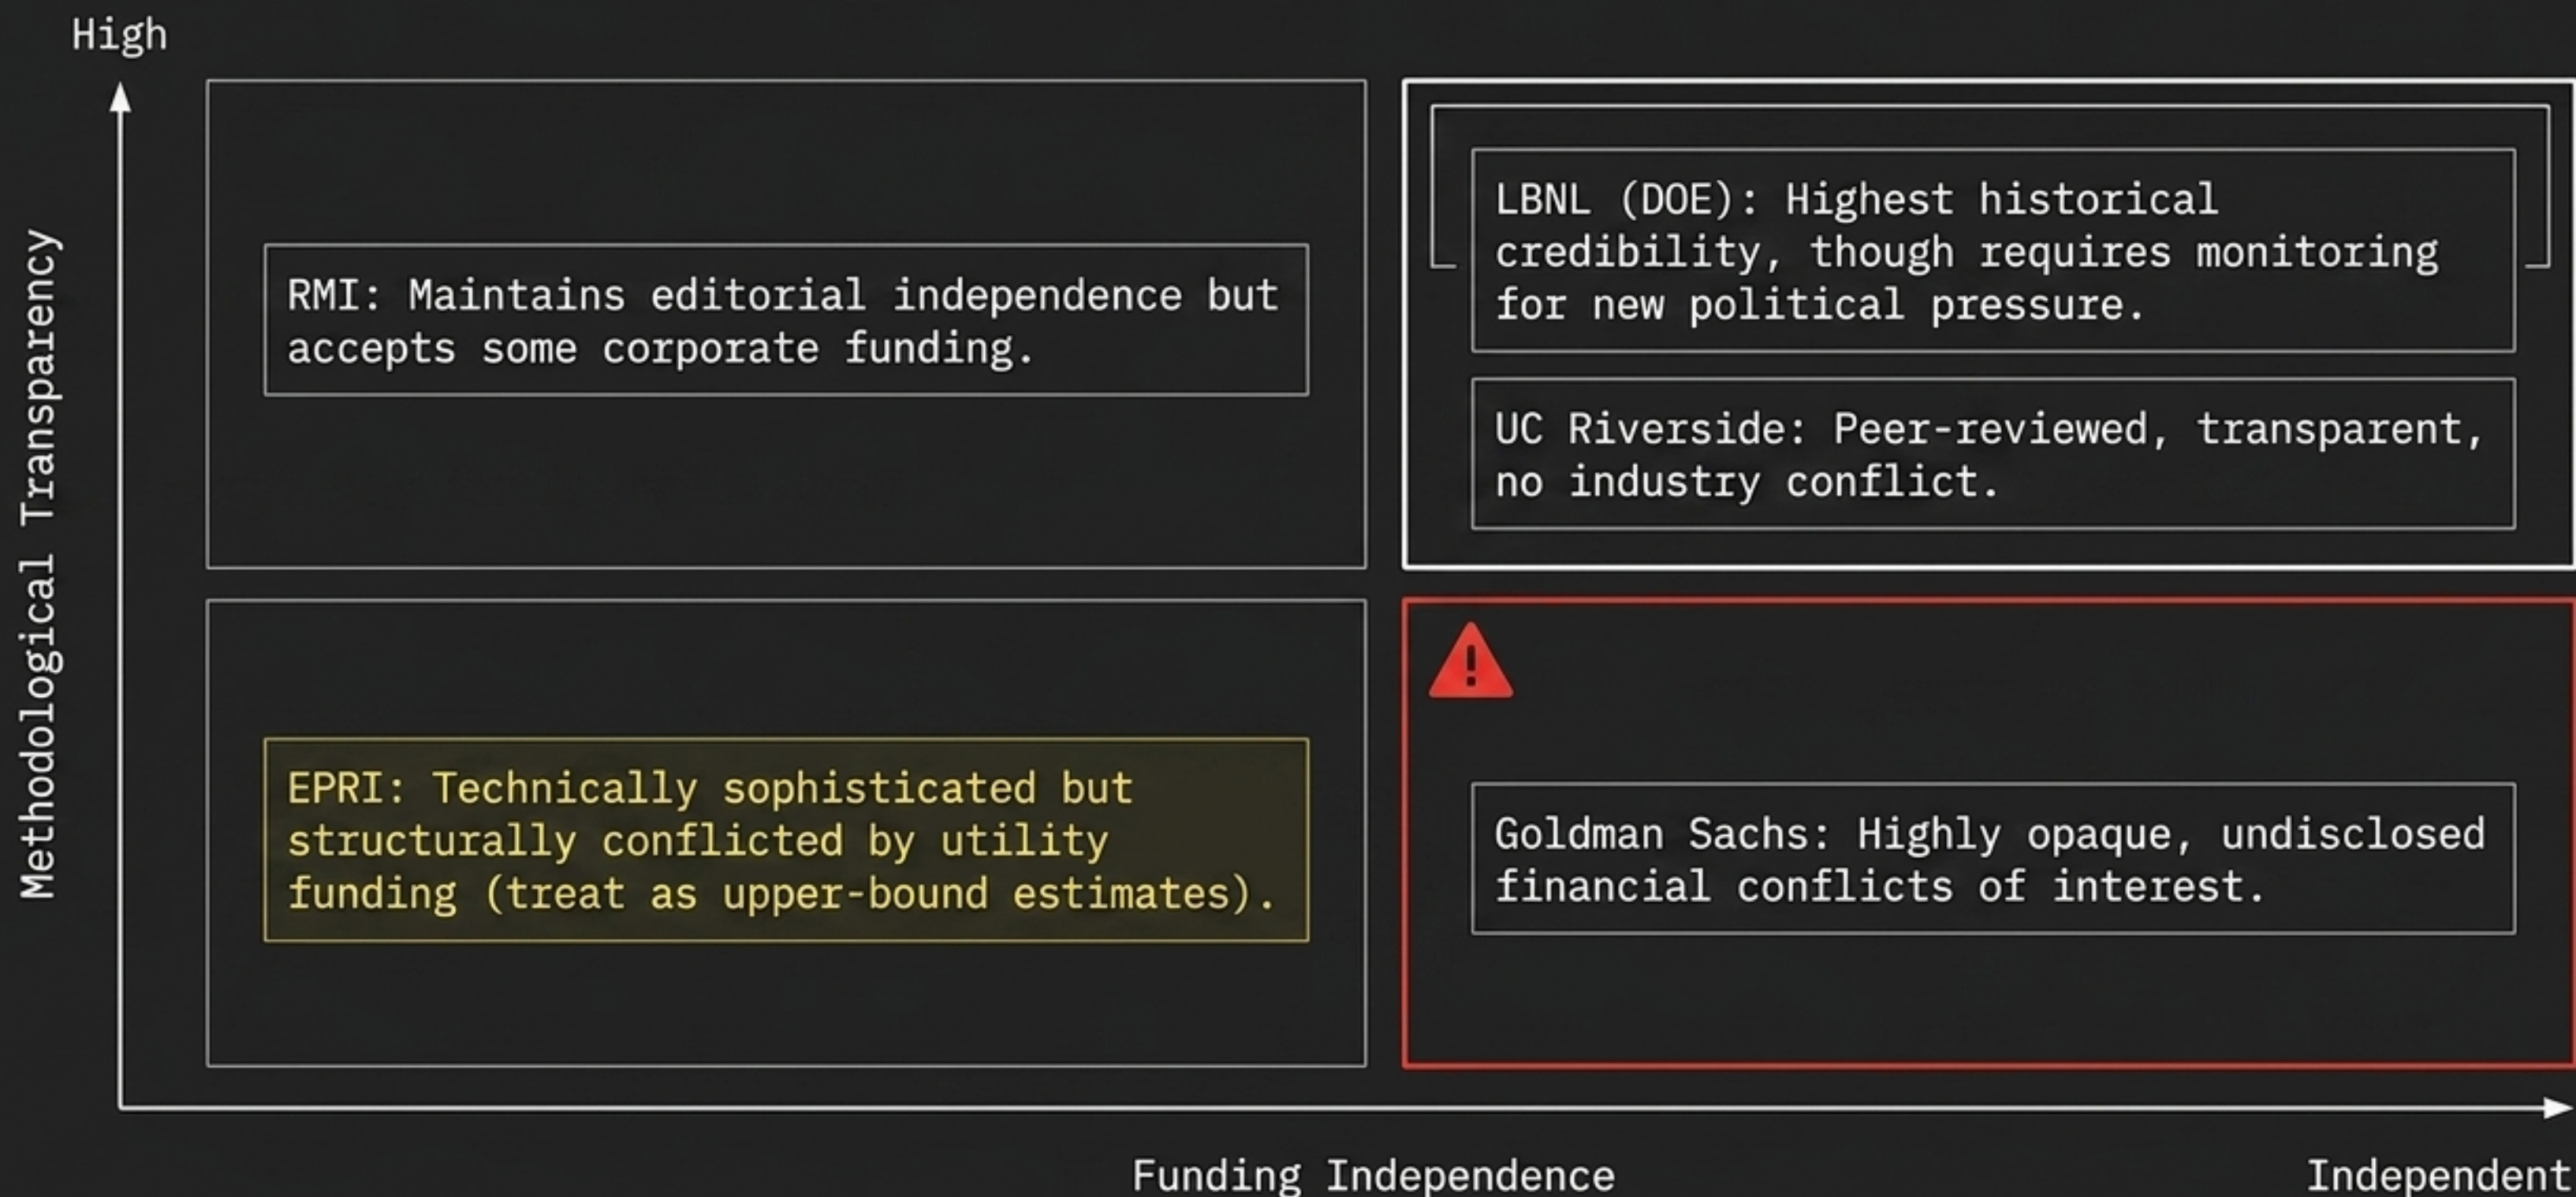
Geographic Redistribution

Will power constraints and water scarcity force a massive geographical shift to regions with excess grid/water capacity?

The Edge vs. The Core

The unknown ratio of future AI inference workloads that will be processed on-device (phones/laptops) versus inside centralized mega-facilities.

The Trust Framework for Future Research



Closing the Divide

The true environmental and social cost of AI infrastructure is being systematically obscured by cherry-picked metrics, undisclosed funding, and political interference at the highest levels of federal regulation.

- 1 . Mandate Comprehensive Disclosure**
Require Scope 3 water and power reporting for all AI data centers.
- 2 . Enforce Funding Transparency**
Mandate conflict-of-interest disclosures in all academic and economic research used for public permitting.
- 3 . Restore Evaluation Independence**
Reverse the political restructuring of NEPA and reinstate objective EPA environmental evaluation procedures.